

The Role of PASTA in Network Measurement

François Baccelli
INRIA-ENS, Ecole Normale Supérieure, France
Francois.Baccelli@ens.fr

Sridhar Machiraju*
Sprint ATL, California, U.S.A
Machiraju@sprint.com

Darryl Veitch†
Dept. of E&E Engineering
University of Melbourne, Australia
dveitch@unimelb.edu.au

Jean Bolot
Sprint ATL, California, U.S.A
Bolot@sprint.com

ABSTRACT

Poisson Arrivals See Time Averages (PASTA) is a well known property applicable to many stochastic systems. In active probing, PASTA is invoked to justify the sending of probe packets (or trains) at Poisson times in a variety of contexts. However, due to the diversity of aims and analysis techniques used in active probing, the benefits of Poisson based measurement, and the utility and role of PASTA, are unclear. Using a combination of rigorous results and carefully constructed examples and counter-examples, we map out the issues involved, and argue that PASTA is of very limited use in active probing. In particular, Poisson probes are not unique in their ability to sample without bias. Furthermore, PASTA ignores the issue of estimation variance, and the central need for an inversion phase to estimate the quantity of interest based on what is directly observable. We give concrete examples of when Poisson probes should not be used, and explain why, and offer initial guidelines on suitable alternative sending processes.

Categories and Subject Descriptors

C.2.3 [Computer-Communication Networks]: Network management; C.4 [Performance of Systems]: Measurement techniques.

General Terms

Management, Measurement, Performance.

Keywords

Active Measurement, PASTA, Probing, NIMASTA.

*When this work was done, the author was also a student at UC, Berkeley.

†ARC Special Research Centre on Ultra-Broadband Information Networks, CUBIN is an affiliated program of National ICT Australia (NICTA)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'06, September 11–15, 2006, Pisa, Italy.
Copyright 2006 ACM 1-59593-308-5/06/0009 ...\$5.00.

1. INTRODUCTION

Poisson Arrivals See Time Averages, or ‘PASTA’, is a property applicable to many stochastic systems. In essence, it states that observations made of a system at time instants obeying a Poisson process, when averaged, converge to give the ‘true’ value, that is to the average that an ideal observer would make when monitoring the system continuously over time. PASTA was first formalized by probabilists, notably in the 1970’s. Wolff, in his classic 1982 paper [22], unified and extended the then-existing PASTA results. The generality of his formulation, based on the ‘Lack of Anticipation Assumption’, which requires simply that the past history of the system does not influence the arrival times of future observers, did away with the need to prove ergodic theorems for each new application, and led to PASTA being widely used.

PASTA has been used [13, 14, 20, 23] to justify the sending of probes (or probe trains) at Poisson epochs in an effort to obtain unbiased estimates of quantities of interest, for example end-to-end delay. However, despite the generality of the PASTA result of Wolff, in many respects the role and utility of PASTA for active probing has become unclear both in the theoretical and practical senses. This paper aims to clarify what Poisson probing, and PASTA itself, can and cannot provide for active probing. In this context, key questions include

- When is PASTA valid in the strict sense?
- When and in what sense is PASTA useful when it holds? is Poisson probing necessarily optimal?
- Are there cases when Poisson probes should *not* be used?
- What role is played by PASTA within the inference problems of active probing?

Related to this last point there is an important, prior question: **What does PASTA apply to?** In other words, Poisson Arrivals See Time Averages, but of what? does PASTA hold for *any* quantity that may form the object of active probing?

Our main focus in this paper is on end-to-end delay over a tandem queueing network, to which PASTA can in fact apply. Delay is a simple, yet important target of active probing measurement in its own right. A natural aim in this context would be to accurately determine any desired statistic of the delay that would be experienced by a single packet of any given size sent into the network in its steady state regime, for example the distribution of such a delay. A particular case is the *virtual work* of queueing theory, which corresponds to the delay a zero sized packet would see, under FIFO scheduling, when sent into the network in steady state. By carefully distinguishing between the non-intrusive case (virtual probes

of zero size) and the intrusive case (real probes of finite size), we provide important insights into the above questions. The simplicity of delay allows rigorous results to be derived, and yet it provides a context rich enough to inform active probing techniques in general.

Our findings group naturally under three distinct categories, and can be summarised as follows.

Sampling Bias versus Intrusiveness

- PASTA states that Poisson sampling is unbiased. In the non-intrusive case, we show that this is not unique to Poisson but is shared by a large class of other sampling processes.
- PASTA states that Poisson sampling remains unbiased even when observers are not virtual, but contribute to system load. Apart from a few exceptions ([11]), this property is not shared by other sampling processes. We argue that it does not follow that Poisson is superior, because of variance and inversion issues, described next.
- We show that *rare probes* can be used to avoid issues of intrusiveness and inversion, and that in such a scenario, Poisson is no longer special anyway.

Bias versus Variance

- PASTA is a statement about bias. It is silent on variance, which is nonetheless of equal importance to estimation.
- There is no general result stating the optimality of Poisson observations with respect to variance or Mean Square Error (MSE), except asymptotically for MSE in the intrusive case¹. Indeed, optimality would in general require a probing stream which is well matched in some sense to network characteristics. In Section 2.2 we give explicit examples showing that Poisson probing can be sub-optimal.

Sampling versus Inversion

- To obtain ‘what one wants’ from what has been observed, for example what the delay distribution would have been if the observers were non-intrusive, based on measurements that were free of sampling bias but which were intrusive, an additional *inversion* step is required. Inversion is typically complex, and in general impacts both bias and variance.
- PASTA is silent on inversion. There is no result stating that Poisson sampling is unbiased, or otherwise optimal, for the full problem of sampling followed by inversion. Furthermore, the zero sampling bias of Poisson in the intrusive case is not necessarily an advantage when it assists in measuring the wrong quantity. It may even be that inversion is impossible, in which case Poisson sampling cannot magically provide unbiased estimates.

Finally, we also strongly emphasize the fact that PASTA does not always hold, as it, in common with alternative probing strategies, requires important conditions to be satisfied.

The picture that emerges is that PASTA plays only a very restricted role in active probing. In a nutshell, active probing estimation seeks to optimise total bias as well as variance performance, and must therefore address both sampling and inversion issues. PASTA deals only with sampling of the available observable, not with inversion to the final quantity of interest, is ignorant of variance, and furthermore excludes the low variance potential of alternative schemes which also enjoy zero sampling bias. In contrast, it’s strength, a lack of sampling bias even in the intrusive case, is not necessarily relevant given the near universal need for inversion.

¹See the end of Section 2.2, and recall $MSE = \text{bias}^2 + \text{variance}$.

We conclude that Poisson probing is poorly motivated, and recommend a *Probe Pattern Separation Rule* as an alternative default for both probes and probe patterns, which offers several advantages.

The remainder of the paper is structured as follows. Section 2 uses simple queueing systems to map out the key issues and concepts, and to give examples of many of the main results and insights in a simple setting. Section 3 introduces the mathematical machinery used to prove results on the bias of sampling processes in a more general setting. We restrict ourselves to hypothetical zero sized probes in this section. In Section 4, we consider probes of non-zero size, examine PASTA, and give a result on *rare probing*, a generic strategy to avoid bias. We then combine our findings and motivate the introduction of a *Probe Pattern Separation Rule* as an alternative to the exponential separation of Poisson probing. In Section 5 we survey selected prior work on PASTA and its application in network measurement, and conclude in section 6.

2. PASTA AND DELAY: THE ISSUES

In this section we illustrate the key facts and issues involved in measuring end-to-end delay from probes, and the role of PASTA, in the simple context of a single, FIFO queue, fed by probe traffic and cross-traffic obeying simple models. As we proceed, we highlight those results which will receive a more general and rigorous treatment in Section 3. In other cases, the results are in the form of counter-examples which prove general points as well as illustrate them.

As several of the examples employ the M/M/1 queue, we summarize some relevant properties here. In the M/M/1 system, packets arrive as a Poisson process of rate λ , and each takes an exponential amount of time, with average μ , to be serviced². To ensure stability and (strict) stationarity of the system, we require the system utilisation $\rho = \lambda\mu$ to satisfy $\rho < 1$. It turns out ([8], pp.202) that the time a packet spends in the system, which is nothing other than its end-to-end delay, is also exponentially distributed with parameter $\bar{d} = \mu/(1 - \rho)$:

$$F_D(d) = P(D \leq d) = 1 - e^{-d/\bar{d}}, \quad d \geq 0, \rho < 1 \quad (1)$$

with mean $\mathbb{E}[D] = \bar{d}$, where D is the random variable representing the delay of a given packet.

A related but distinct quantity is the waiting time W of a packet, which also corresponds to the *delay experienced by a virtual observer of zero size*. This distribution:

$$F_W(y) = P(W \leq y) = 1 - \rho e^{-y/\bar{d}}, \quad y \geq 0, \rho < 1, \quad (2)$$

with mean $\mathbb{E}[W] = \rho\bar{d}$, has an atom at the origin corresponding to the probability $1 - \rho$ of finding the system empty, resulting in zero waiting time, and also zero delay in the case of a zero sized packet.

As mentioned in the introduction, an important quantity in queueing theory is the *virtual work*. This is a stochastic process $W(t)$, defined for continuous time $t \in \mathbb{R}$, which corresponds to the waiting time a packet of size $x = 0$ would experience when sent at time t into the network in steady state. Because this is also the delay for such a zero sized observer, we refer to this here as the *virtual delay*. This process represents the ground truth of the delay (of zero sized packets) in the system. The waiting time distribution W of Equation (2) is simply the marginal of $W(t)$.

The other queues we consider in this section are studied using Monte Carlo simulation implemented in C and Matlab. The queue ‘simulation’ directly implements the Lindley recursion (see

²Often μ is taken to be the service *rate* instead.

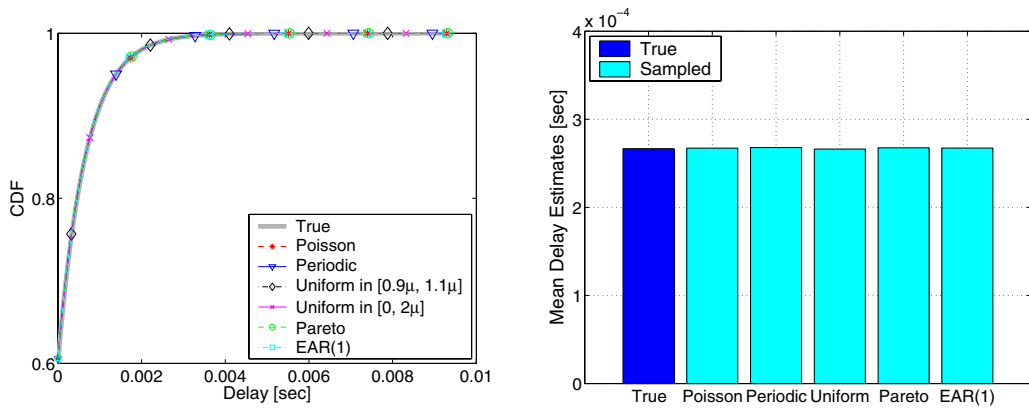


Figure 1: Sampling bias of delay, non-intrusive case ($x = 0$). Left: CDF as seen by various probing streams, and the true delay distribution. Right: resulting mean estimates. Each probing stream is unbiased.




Figure 2: Sampling bias of delay, intrusive case ($x > 0$). Left: CDF as seen by various probing streams, and true delay distributions (one per stream, the closest grey curve in each case). Right: resulting probe based mean estimates, and true means. Each probing stream results in a new true delay distribution, which is sampled with bias by the probes, except the Poisson case (PASTA).

eg. [8]) on waiting times defining the system and is exact to machine precision. Two kinds of statistics are collected. First, per-packet delay values, from which the delays, regardless of packet size, due to probe traffic or cross-traffic can be extracted. Second, the waiting time distribution W is obtained by observing the virtual delay process $W(t)$ continuously over time. As it is stored in histogram form, there is a discretisation error. However, this error can be bounded, and we control it in each case so that errors are negligible on the scale of the plots given. Similarly, we use long simulations of 1000000 probes to make confidence intervals small or negligible (in the latter case we don't show them), and employ warmup periods of at least $10\bar{d}$ to damp transients. From the observed $W(t)$, we obtain the distribution of D for non-zero probes by convolving with the probe size distribution.

Finally, recall that if a is some parameter of stochastic data, then an estimator of a is any function \hat{A} of the data designed to estimate a ; this estimator is unbiased if $a = \mathbb{E}[\hat{A}]$.

2.1 Bias

We will consider each of sampling bias in the intrusive case, sampling bias in the non-intrusive case, and inversion bias. Five different arrival processes - including 'Poisson', 'Uniform', 'Pareto', 'Periodic', and 'EAR(1)', will be used for probes, in order to offer a spectrum of bursty behaviors. Three of these, namely 'Poisson', 'Uniform', and 'Pareto', are *renewal processes*, that is inter-arrivals

are given by independent and identically distributed (i.i.d.) random variables. They differ only in the choice of distribution, respectively exponential (yielding a Poisson process), uniform, and the heavy tailed Pareto (with finite mean but infinite variance). The 'Periodic' probe stream is also renewal but in a very degenerate sense as inter-arrivals are constant, and is best regarded as a deterministic stream (a random uniform phase, determining the offset of the periodic grid from the time origin, makes it stationary despite this rigidity). The EAR(1) process, described in detail below, has correlated inter-arrivals with exponential marginal. For the sake of brevity, details such as complete parameter settings will be omitted. Figure 1 gives results for each of the above probing streams, with a shared average inter-probe spacing, using probes of zero size. Consequently, there is no issue of intrusiveness - probes do not affect the system, nor of inversion - we are directly measuring what we wish to measure. The issue of sampling bias can therefore be seen in isolation.

The grey curve in the plot on the left of Figure 1 shows the true Cumulative Distribution Function (CDF) for the delay of zero sized probes, calculated from Equation (2). As expected, the curve corresponding to Poisson probes agrees with the true one: as is well known, PASTA applies to this system. However, each of the other 5 curves overlay the true result equally closely. In this non-intrusive case, the lack of sampling bias of Poisson probing is shared by many other probing schemes. The (tightly estimated) expected de-

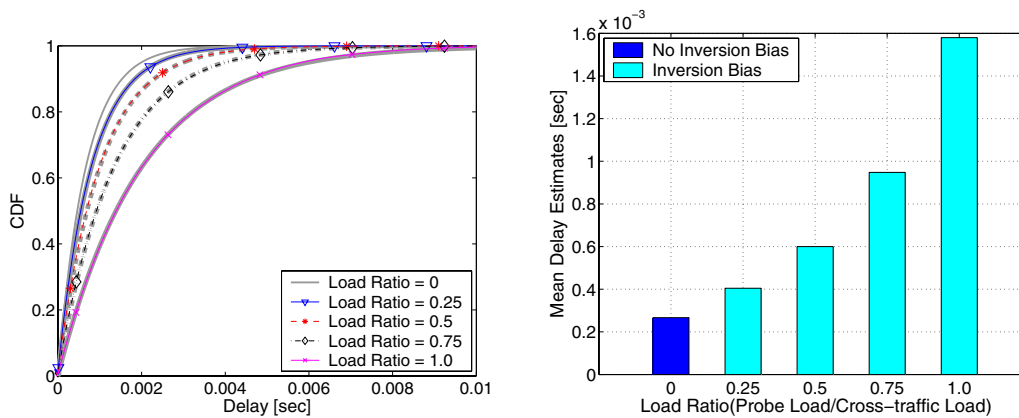


Figure 3: Inversion bias of delay, range of intrusiveness ($x \geq 0$). Left: CDF as seen by Poisson probing streams of different rate, and true delay distributions (one per stream), as well as the true unperturbed delay (with no probes). Right: corresponding mean delays as a function of probe to total load ratio. PASTA eliminates sampling bias, but total system behavior increasingly deviates from that of the unperturbed system.

lays in the right plot of the figure confirm this by agreeing with the true value in each case. In Section 3, our main result is to prove that a wide class of processes share this desirable property, and for far more general systems than the simple M/M/1 queue.

In Figure 2 we consider the same probing arrival streams, but allow the probes to have a service time $x > 0$ (for simplicity x is a constant, but this is not essential). As a result, intrusiveness becomes an issue: probes do affect the system, both in load and in more detailed characteristics. To avoid dealing with inversion issues, for the moment our objective is to know the true delay of the full system, combining cross-traffic and non-virtual probes, that a packet of service time x would experience. In other words, again we seek to measure the same object that we have direct access to through probing.

The mean estimates in the right hand plot of Figure 2 confirm that each probing stream results in a different system behavior (despite equal loads), and shows that each now gives a biased estimate of its respective $E[D]$, with the exception of ‘Poisson’. Hence, PASTA continues to hold, whereas the other probing streams, despite being unbiased when $x = 0$, now suffer from a bias due to intrusiveness. The corresponding CDFs in the left hand plot show in greater detail how the systems are different for each stream, and how the bias varies as a function of delay. These results illustrate that PASTA holds in the intrusive case. In Section 4, we will state the general conditions under which PASTA can be expected to hold.

We now study inversion bias in isolation. We achieve this by employing Poisson streams exclusively, thereby benefiting from their zero sampling bias in all cases. Furthermore, we let the probe service time X be exponentially distributed with the same parameter, μ_T , as for the cross-traffic packets. This results in a probe+traffic system which is still M/M/1, with rate $\lambda = \lambda_T + \lambda_P$ and average service time μ_T , enabling Equation (1) to be used.

Figure 3 shows the unsurprising but significant fact that increasing the probing load through increasing λ_P results in the overall system deviating increasingly far from the original unperturbed system in which $\lambda_P = 0$. Consequently, even if an estimate of the true mean (or the CDF) is unbiased, that estimate is an estimate of the **full** (probe + cross-traffic) system, **not** the unperturbed (cross-traffic only) system that one wishes to measure. Thus, ‘what we want is not what we directly measure’. To obtain the unperturbed delay from the perturbed one requires an entirely separate inversion step, which, even in this very simple example of inverting one

kind of delay to another, is highly non-trivial except for the simplest one hop models. For other inference objectives common in active probing, such as using packet-pair methods to estimate bottleneck bandwidth, the degree of inversion required, and therefore its potential impact, is far greater. Another way of seeing this is to note that probes sent as a Poisson process at the sender will not arrive as Poisson process at the bottleneck link in general, and will also be affected by their onward passage from that link to the receiver. Thus, the probes are ‘sampling’ the bottleneck link, but not in a Poisson way, and not in isolation.

We have two contributions to make on the inversion issue. First, we point out that not only can it be extremely challenging, but more importantly that it is a difficulty which arises generically for which PASTA offers no solution. Indeed, it was recently shown in [10] that, even in a simple 1-hop system, unless the cross-traffic obeys particular restrictive assumptions, full knowledge (i.e. the law) of the cross-traffic process feeding the hop is unobservable. In such a case, strict inversion is impossible even in principle. This corresponds to an extreme form of the inversion problem, relating to fundamental issues of system non-identifiability, which PASTA is powerless to mitigate. Second, in spite of the difficulties described above, there is a general way, the *rare probing* of Theorem 4, to address inversion bias in the intrusive case, which we explore in Section 4.

2.2 Bias versus Variance

In Section 2.1 we focused on the bias of estimators based on a simple average of delays experienced by probes. In this section we look into the variance of these estimators. When bias is non-zero, we examine bias-variance tradeoffs in the traditional manner by considering the mean squared error $MSE = \text{bias}^2 + \text{variance}$.

Thus far we have considered cross-traffic packets arriving as a Poisson process. However, in general cross-traffic will interact with probe traffic in ways which depend on the correlation or ‘burstiness’ structure of each, and estimation variance will be a function of these interactions. To show this clearly we need a richer context than the simple memoryless structure of the Poisson process. We use the *Exponential (first order) Autoregressive process* (EAR(1)) [5] as a convenient way to generate a point process for CT arrivals which has a well defined correlation time scale, that is a scale τ^* below which autocorrelation is significant, but above which it decays rapidly to zero. Like the Poisson process, it consists of exponen-

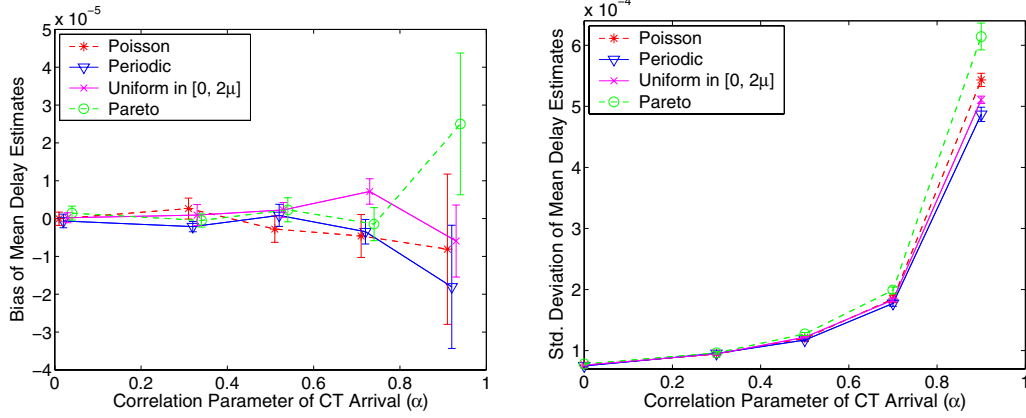


Figure 4: Bias and variance of delay with correlated cross-traffic, non-intrusive case ($x = 0$). Left: bias of mean estimates seen by different probing streams as a function of the EAR(1) parameter α of the cross-traffic, using 100000 probes. Right: corresponding estimates of standard deviation. Although all probing schemes are unbiased, their variances differ, and Poisson is not the smallest.




Figure 5: Bias, variance and MSE of delay with correlated cross-traffic, intrusive case ($x > 0$). Left: bias of mean estimates seen by different probing streams for $\alpha = 0.9$ as a function of intrusiveness probe load/total load. Middle: corresponding estimates of standard deviation. Right: corresponding $\sqrt{\text{MSE}} = \sqrt{\text{bias}^2 + \text{variance}}$. Only the Poisson probing is unbiased, but the scheme with minimal MSE depends on α .

tial inter-arrivals of intensity λ , but unlike it, inter-arrivals form a positively auto-correlated AR(1) process, with correlation structure

$$\text{Corr}(i, i + j) = \alpha^j, \quad j = 0, 1, 2, \dots, \quad \alpha \in [0, 1). \quad (3)$$

The delay constant $j^*(\alpha)$ of this geometric decay (extracted by setting $\alpha^j = e^{-j/j^*}$) translates to a correlation time scale of $\tau^*(\alpha) = j^*/\lambda = (\lambda \ln(1/\alpha))^{-1}$, which rises from 0 when $\alpha = 0$ (the Poisson case), to ∞ as $\alpha \rightarrow 1$.

Figure 4 shows the effect of increasing α on the estimation of mean delay, for four different non-intrusive probe streams of identical rate. In the left plot we see, as expected, a lack of bias for each stream (note the vertical scale and confidence intervals, offset for visibility), in agreement with the results of Figure 1, regardless of the value of α . In contrast, the right hand plot shows that the standard deviation of the estimates separate at large α . The important point to note is that this separation clearly exceeds the confidence intervals: the Poisson stream has higher variance than either Periodic or Uniform. This is a counter-example making the general point that Poisson sampling does not imply minimal variance.

We now offer some insight into why Poisson probing gives rise to higher variance than Periodic in this case. First note that as the correlations in the cross-traffic increase, so do those of the virtual delay process $W(t)$ itself. If we could make estimates based on

$W(t)$, they would therefore have increased variance³. Each probing scheme samples $W(t)$, experiences the larger correlation at higher α , and thereby inherits the larger variance. Exactly how much variance however depends strongly on the details of the sampling scheme, not merely on the average sampling rate λ_P . Bearing in mind that samples which are closer together will be more correlated, periodic probing has the advantage of guaranteeing a minimum distance between them. It can therefore ‘jump over’ correlation inducing bursts, provided that $1/\lambda_P$ is large compared with the correlation scale of $W(t)$. In contrast, in a Poisson process, arrivals may be much closer than $1/\lambda_P$ with appreciable probability, increasing the correlation considerably between such samples. In the example here, $1/\lambda_P \approx 20\tau^*$ even for $\alpha = 0.9$, so the periodic stream produces close to i.i.d. samples in all cases.

In Figure 5 we consider the intrusive case for a wider range of candidate probing schemes. We fix α at 0.9, and examine dependence as intrusiveness is increased by increasing probe size, shown as a function of the ratio of probing load to CT load. The left-most plot shows that bias is now present, and increases with α , for all schemes except for Poisson (PASTA). The variance results of the middle plot echo those seen in Figure 4: there are schemes

³It is well known [2] that the variance of the sample mean calculated over a time window of given width is essentially the integral of the correlation function over the corresponding range of lags.

which perform both better and worse than Poisson. The rightmost plot in Figure 5 combines bias and variance, and we see the trade-offs at work: the relative overall performance of different schemes changes with x . In particular, as bias becomes stronger for its competitors at load ratios above 0.12, Poisson begins to outperform Periodic, but continues to be outdone by the Uniform renewal with wide support.

In this section we have presented only a few illustrations of what is a general point: PASTA is silent on estimation variance, and the performance of Poisson probing, for general cross-traffic processes, plays no privileged role with regard to variance. The same holds true for MSE, with one exception. Asymptotically, as the number of samples tends to infinity, the variance of any consistent estimator will tend to zero, resulting in the asymptotic MSE being equal to the bias squared. In the intrusive case, this clearly gives the advantage to Poisson probing. In general however, overall statistical performance is a function of how well the probing stream is adapted to the cross-traffic, and the nature of that traffic. The optimal approach (if any) will also be strongly determined by the choice of constraints such as measurement duration, probe budget, and acceptable intrusiveness profile.

2.3 The Need for Technical Assumptions

Thus far in this section we have passed over the issue of technical assumptions. For instance, we assumed that PASTA holds universally. In fact, its validity is dependent on other conditions, most importantly the *lack of anticipation assumption (LAA)* mentioned in the introduction. Similarly, while Figure 2 indicated that non-Poisson probing schemes have zero bias, this is true only if suitable *joint ergodicity* conditions on the cross-traffic and probe traffic are satisfied. Intuitively, ergodic systems are those which are in some sense free enough to explore, in an unbiased way and on a single sample path, the full range of behavior which one would find if one could examine *all* sample paths. It is also important to note that zero bias, although an important objective, is not enough. To be useful in practice, we also need measurements to converge to the parameter to be estimated on a *single* sample path, as we witnessed in each example thus far. To ensure this, again suitable ergodic conditions are required. We defer a rigorous description of these issues until the next two sections, and complete this one with an illustrative example and intuitive explanation of these issues.

Figure 6 gives the outcome of a non-intrusive experiment which is identical to that reported in Figure 1, but with one crucial difference: the Poisson arrival times of cross-traffic have been replaced by periodic arrivals of the same average intensity (packet sizes have not been altered). Each probing stream continues to measure the mean delay, and even the entire delay distribution, without bias, with the exception of the periodic probe stream, which is markedly different. In fact, since the period of the Periodic stream is equal to an integer multiple of the cross-traffic period (equal to 10 in this case), the two streams are effectively ‘phase locked’, and in such a case the joint ergodicity conditions are not satisfied. As a result, the probes can never sample average conditions on this sample path alone, but only those found at a particular point in the cycle of the cross-traffic arrivals.

Despite the rigidity of the periodic cross-traffic, the other probing streams do satisfy the required joint conditions since they are each *mixing* processes, which is a stronger form of ergodicity. Similarly, the joint ergodicity assumptions were satisfied in the cases shown earlier of periodic probe traffic and either Poisson or EAR(1) cross-traffic, since these latter processes were mixing. In a loose sense, they provided enough variability to overcome the rigidity of the periodic probes. We discuss this in more detail in the next section.

3. NON-INTRUSIVE MEASUREMENT

In the previous section we discussed sampling bias, inversion bias, variance, and contrasted the intrusive and non-intrusive cases in parallel. In this section we focus on sampling bias in the non-intrusive case only, leaving the intrusive case (and most of our comments on inversion) to Section 4. Our aim is to expand in a rigorous way on the observations of Figure 1, that many processes other than Poisson enjoy zero sampling bias in the non-intrusive case. Although this case may not seem to be practically useful, since real probes must always intrude, it turns out that the design of effective practical probing is based on a clear understanding of this case. The basis of our treatment is the machinery of ergodic theory [16] and Palm calculus [1], in the context of marked point processes [3, 1]. We give an overview of these areas before proving our main result, a description of precisely when the zero bias property is true.

3.1 Setting

We adopt a setting which allows for very general probe traffic, cross traffic, and network models. We make the following assumptions: *stationarity* of the probe traffic, cross-traffic and the network behavior, *ergodicity* of the cross-traffic and probe traffic, and *independence* of the probe traffic from the cross-traffic. The assumptions on probes are not restrictive since the experimenter controls the probing stream.

We model probe traffic as a (strictly) *stationary point process* \mathcal{P} of intensity λ_P . That is, a sample path of the process is simply the set of times $\{T_n\}$ at which the (zero sized) probes arrive, and there exists a probability law $P_{\mathcal{P}}$ that determines the probability of any event concerning sample paths. For example, it governs the probability that the first inter-arrival time after the time origin: $T_2 - T_1$, equals or exceeds the mean value $1/\lambda_P$ (this would be 1 for periodic probes or e^{-1} for a Poisson stream), as well as the probability that n probes fall in a time interval I_1 and m in I_2 , for intervals I_1 and I_2 . No constraints are placed on $P_{\mathcal{P}}$; we allow any structure of probe arrivals (provided points do not coincide).

We model cross-traffic as a *marked stationary point process* \mathcal{T} of intensity λ_T . As before this consists of the arrival times of packets, but now also *marks*, random variables associated with each packet which give additional information about the traffic. This includes the random packet size, but also anything else which characterises the stream, for example the packet sizes may depend on the arrival patterns, or packet sizes, of previous arrivals. The probability law $P_{\mathcal{T}}$ governs all details of \mathcal{T} , both of arrival times and marks.

The model of an end-to-end path typically used in active probing is essentially the *tandem queueing network* of queueing theory. It consists of a set of FIFO queues and transmission links in series, each with its own independent cross-traffic stream. Packets from a given stream are all *n-hop-persistent* (traversing n hops before exiting) and frequently n is simply taken to be 1 for each stream. Our network setting is more than general enough to cover such a model. However, it is not explicitly defined in terms of queues like the one above, but instead operates in an abstract setting, which allows for even greater generality. It is capable of including cross-traffic streams correlated across nodes, cross-traffic with feedback such as TCP, non-FIFO scheduling disciplines, varying over nodes, probes which follow different paths through a network (modelling load balancing), and more, provided the above technical assumptions are satisfied. Technically, each of these cross-traffics, and their dependencies, are contained in a single marked point process \mathcal{T} , where the marks carry most of the detailed information, such as which nodes are traversed by a given packet. In this way, much of our general network model, in fact all of its stochastic components, is subsumed into a rich cross-traffic description.




Figure 6: Sampling bias of delay with non-mixing cross-traffic, non-intrusive case ($x = 0$). Left: CDF as seen by various probing streams, and the true delay distribution. Right: resulting mean estimates. Each probing stream is unbiased, except for periodic.

The details of the queueing itself are not contained in \mathcal{T} but would have to be specified separately, for example, if one wished to simulate the network. However, provided everything that is not in \mathcal{T} acts deterministically on the cross-traffic and probe inputs, it does not need to be specified explicitly. For example, our results hold ‘for free’ for each of FIFO, Weighted Fair Queueing, or processor sharing queueing disciplines, since each of these is deterministic given the traffic inputs.

The final components of the basic setting is to specify the *observable*, which is the quantity related to probes that we have access to, and the *ground truth*, that is the quantity we wish to estimate. In the case of active probing the available data is simply the arrival times of probes to the receiver, or equivalently (since the sending times T_n are known), their end-to-end delays. Since we consider only the non-intrusive case in this section, the underlying observable or ‘ground truth’, which we denote by $Z(t)$, $t \in \mathbb{R}$, will be taken to be the virtual delay process $W(t)$ from Section 2. Recall that this is the delay experienced by a zero sized observer entering at time t .

Our main goal is to learn about the process $Z(t)$. Technically, this reduces to determining the expectation $E[f(Z(t))]$ of some positive function f of $Z(t)$. The choice of f gives us great freedom in the kind of statistic we may wish to measure. Good examples, which we have already met in Section 2, are the identity (giving us the mean delay) or an indicator function noting whether $Z(t)$ is smaller than some threshold (giving us the CDF of delay). More will be considered below.

With the setting established, we now indicate where sampling and inversion fit in to it. Probes sent at times $\{T_i\}$ literally sample $Z(t)$ at those times. Hence, the values $f(Z(T_1))$, $f(Z(T_2))$, \dots , are what is available to estimate $E[f(Z(t))]$. As we send more probes, we have more samples and expect our estimates to improve. Specifically, we want the following *almost sure convergence*:

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N f(Z(T_n)) = E[f(Z(0))]. \quad (4)$$

For instance, if f is the identity function, the right hand side is the mean virtual delay to which the sample mean estimate on the left hand side must converge. Note that stationarity implies that $E[f(Z(0))] = E[f(Z(t))]$ for any time t .

3.2 Ergodic Theory and Palm Calculus

Statements like Equation (4), where an empirical average (the left hand side) is equivalent to an ‘ensemble mean’ or mathematical expectation (the right hand side) are known collectively as er-

godic theorems [16]. Intuitively, this equivalence arises because a single sample path of an ergodic process will over time come to resemble every other sample path, with more extreme paths taking appropriately longer to emerge. In Section 3.3 we determine when Equation (4) holds. To do so, we must first introduce key aspects of ergodic theory and Palm calculus.

3.2.1 The Joint Law and the Product Space

To deal with ergodicity of the whole system, we must know the joint law governing both probe and cross-traffic. Because these are independent, the events in the combined system can be described through the *product space* of \mathcal{P} and \mathcal{T} , denoted by \mathcal{F} , with an associated probability law P which is the product of $P_{\mathcal{P}}$ and $P_{\mathcal{T}}$.

Intuitively, (\mathcal{F}, P) enumerates all sample paths of the combined system and their associated probabilities. The following example illustrates this. Consider a system in which probe and cross-traffic are each periodic with a period of 1 time unit (for simplicity we ignore the marks of \mathcal{T}). Each probing sample path is completely determined by its *phase* $y \in [0, 1) = T_1$, the distance from the time origin to the first probe. Similarly, the cross-traffic is described by a phase $z \in [0, 1)$. We take the phases to be uniformly distributed over $(0, 1]$, so that each stream is stationary. Thus, each sample path of the combined system is uniquely described by (y, z) , which is an element of the product space $\mathcal{F} = [0, 1) \times [0, 1)$. Assuming independence between the streams, the joint probability P is just

$$\begin{aligned} P((y, z) \in [a, b] \times [c, d]) &= P_{\mathcal{P}}(y \in [a, b]) \cdot P_{\mathcal{T}}(z \in [c, d]) \\ &= (b - a)(d - c), \end{aligned}$$

where $b \geq a$, $d \geq c$.

3.2.2 Ergodicity and Time Shifts

We will describe the stationarity and ergodicity of a point process by means of the associated *shift operator* $\{\theta_t\}$, which represents a shift in time of value $t \in \mathbb{R}$ of the whole sample path of the process, that is of all of its points. Stationarity is very simply expressed in this framework: by definition, it holds if, for all t , a shift of t leaves the law of the process unchanged. In order to define ergodicity, we have to first introduce the notion of *invariant event*.

Invariant Event: An *invariant event* for a shift operator $\{\theta_t\}$ (for instance that of a point process) is an event A such that $A = \theta_{-t}(A)$ for all $t \in \mathbb{R}$. An example of such an event is a set of paths each of which has an infinite number of inter-arrivals larger than some value x , because translation would not change this property for any path in A , so $\theta_{-t}(A)$ would contain exactly the same paths.

Ergodic Shift: A shift is *ergodic* if all of its invariant events are trivial; i.e. of probability either 0 or 1. For instance, all stationary renewal processes are ergodic (e.g., see Section 1.6 of Chapter 1 in [1]). In particular, a periodic point process with a random phase that is uniform over the period is both stationary and ergodic.

Let us return to the product space defined above. One can define a *product shift* which operates, simultaneously but independently, on both \mathcal{P} and \mathcal{T} . We would like this product shift to be ergodic, in which case we say that \mathcal{P} and \mathcal{T} are *jointly ergodic*. However, a pitfall here is that for this to be true, it is not enough that one or the other, or even both, of the probe process and cross-traffic process be ergodic in their own right. To see the significance of this, we continue the ‘periodic-periodic’ example from Section 3.2.1.

Let the event A be those sample paths where, for all n , $T_n - C_n < 0.25$, where T_n and C_n are respectively the arrival time of the n -th probe and n -th cross-traffic packet. For our phase locked example, this translates to $P(y - z < 0.25) = 0.25$. However since the offset between the two streams is fixed at $y - z$ for all n , A is an invariant event, yet it has probability which is neither 0 nor 1. Hence, the product shift is not ergodic, despite each of the streams being individually ergodic. This example is the simplest one of its kind, but is by no means the only one. Such periodic behaviors are actually common in IP networks, for example when dealing with a small number of persistent TCP flows on an access link.

3.2.3 Palm Probability

As pointed out in Section 3.1, probing fundamentally involves sampling. Whereas the underlying observable $Z(t)$ is a process in continuous time, in practice we must work with discrete observations taken at stochastic times. *Palm Calculus* [1] is a body of results dealing with how the probabilistic laws governing a process in fact depend on the ‘conditional viewpoint’ from which they are ‘observed’.

On our product space, the *Palm probability* \mathcal{P}_0 with respect to the probe point process is defined by

$$\mathcal{P}_0(Z(0) \in B) = \frac{1}{\lambda(b-a)} E \left[\sum_{T_n \in (a,b]} 1_{Z(T_n) \in B} \right],$$

for any real numbers $a \leq b$, where λ is the intensity of the probing point process. Hence, the Palm probability is the average fraction of probes in $(a, b]$ which observe $Z(t)$ as being in the set B . As B is arbitrary, this is equivalent to knowing the entire distribution.

We now show that when probes are non-intrusive and independent of cross-traffic, then $\mathcal{P}_0[Z(0) \in B] = P[Z(0) \in B]$, for all B . Let $N(a, b]$ denote the number of probes in $(a, b]$. From the independence assumption and the fact that Z is a function of the cross-traffic point process only, we have

$$\begin{aligned} \mathcal{P}_0(Z(0) \in B) &= \frac{1}{\lambda(b-a)} E \left[\sum_{T_n \in (a,b]} 1_{Z(T_n) \in B} \right] \\ &= \frac{1}{\lambda(b-a)} \sum_{n \geq 0} P_{\mathcal{P}}(N(a, b] = n) \int_{\mathbb{R}^n} \\ &\quad \sum_{k=1}^n P_{\mathcal{T}}(Z(t_k) \in B) f_n(d(t_1, \dots, t_n)), \end{aligned}$$

with $f_n(d(t_1, \dots, t_n))$ denoting the law of the epochs of the n probes falling in the interval $(a, b]$ given that there are exactly n probes in this interval. By stationarity, for all k $P_{\mathcal{T}}(Z(t_k) \in B) =$

$P_{\mathcal{T}}(Z(0) \in B)$. Hence

$$\begin{aligned} \mathcal{P}_0(Z(0) \in B) &= \frac{1}{\lambda(b-a)} E_{\mathcal{P}} \left[\sum_{T_n \in (a,b]} P_{\mathcal{T}}(Z(0) \in B) \right] \\ &= P(Z(0) \in B) \frac{1}{\lambda(b-a)} E \left[\sum_{T_n \in (a,b]} 1 \right] \\ &= P(Z(0) \in B). \end{aligned} \quad (5)$$

This result implies that the underlying law of the observable as seen by independent probes (represented by \mathcal{P}_0) is equal to the ‘ground truth’ (represented by P).

3.3 NIJEASTA and NIMASTA

In this section we give two theorems which build on the above to establish broad conditions under which Equation (4) holds.

The significance of the joint ergodicity property stems from the *pointwise ergodic theorem* of Birkhoff. This theorem implies (see Section 1.6 in Chapter 1 of [1]) that when the product shift is ergodic, for all stationary stochastic processes $Z(t)$ defined on the product space, and for all positive functions f , the following (almost sure) limit exists

$$\lim_{N \rightarrow \infty} \frac{1}{N} f(Z(T_n)) = E^0[f(Z(0))] \quad (6)$$

where E^0 denotes the Palm probability of the probe point process. We can now give our main result

THEOREM 1. *If the product shift is ergodic and the probing stream is independent of cross-traffic, then Equation (4) holds.*

Proof: Since the product shift is ergodic, it follows from Property 1.6.3 pp.52 in [1] that the discrete shift θ_{T_1} is ergodic w.r.t. the Palm probability \mathcal{P}^0 , so that for all positive functions f the limit

$$\lim_{N \rightarrow \infty} \frac{1}{N} f(Z(T_n)) = E[f(Z(0))]$$

holds almost surely. The result then follows from (5). \square

This result can be summarised as:

NIJEASTA: *Non-Intrusive Jointly Ergodic Arrivals See Time Averages.*

The jointly ergodic assumption of NIJEASTA is similar to the Lack of Anticipation Assumption of Wolff in that it states exactly what is required, but does not say when it is true, which can be inconvenient in practice. Our next theorem is classical ([16], Theorem 6.1, pp.65) and states simple sufficient conditions under which the joint ergodicity holds, based on the idea of *mixing*. A point process \mathcal{P} (or equivalently its shift) is said to be *mixing*⁴ if, for all events A, B :

$$\lim_{t \rightarrow \infty} P_{\mathcal{P}}(A \cap \theta_{-t}(B)) = P_{\mathcal{P}}(A)P_{\mathcal{P}}(B).$$

Intuitively, mixing is a special (and strong) form of ergodicity where on separation under the shift, all memory between any sets A and B is lost, so that they ultimately act as independent events.

THEOREM 2. *The product space \mathcal{F} of \mathcal{P} and \mathcal{T} is ergodic whenever at least one of them is a mixing process, and the other ergodic.*

Of the two cases covered here, that of a mixing probe process has practical importance, because although we may suspect that cross-traffic is mixing, especially in the Internet backbone where myriads of random effects wash out deterministic synchronisation, we cannot guarantee it. On the other hand, if we choose to use probing processes which are mixing, we are assured of satisfying the joint

⁴In fact both *weak* and *strong* mixing can be defined [16]




Figure 7: Simulation illustrating NIMASTA in a multi-hop system, and sampling bias due to phase-locking. Left set of curves: periodic cross-traffic on hop 1, Right: window-constrained TCP flow on hop 1.

ergodicity conditions required for zero sampling bias, regardless of the dynamics of cross-traffic. To highlight this property, which generalises PASTA (in the non-intrusive case!), we coin:

NIMASTA: *Non-Intrusive Mixing Arrivals See Time Averages.*

It is useful to review at this point the observations of Section 2. Three kinds of processes appeared there: stationary renewal processes (with exponential, uniform, or Pareto inter-arrivals), the periodic process (with random phase), and the EAR(1) process. As is well known [3], renewal processes are mixing provided that the support of the inter-arrival distribution contains an interval where the density is larger than a positive constant, and the EAR(1) process is also strongly mixing [5]. However, the periodic process is not, although it is ergodic. The non-intrusive examples throughout Section 2, in particular in Section 2.3, illustrate NIMASTA and NIJEASTA at work, depending on which traffics are mixing or not. Finally, note that it is easy to construct a great variety of mixing processes, for example using Markov processes with a particular structure, as we will see in the next section.

Figure 7 illustrates NIMASTA for a more complex network: a 3 hop FIFO route with capacities of [6, 20, 10] Mbps. Each hop carries 1-hop-persistent cross-traffic which, for hop [1, 2, 3] respectively, were chosen to be [periodic, Pareto, TCP] or [TCP, Pareto, TCP], combinations which include potential for phase locking, long-range dependence, and feedback (a single TCP flow is used). For space reasons we omit many details in this ‘proof of principle’ example, which was generated using ns-2 [21], modified to allow the ground truth $Z(t)$ to be calculated accurately (see Section 7). Two sets of results are given in the figure, depending on whether the cross-traffic on hop 1 is periodic, or window-constrained TCP. In each case the delay marginals show that NIMASTA holds for each of the mixing probe traffics, but not for the periodic probes, which become phase locked. In the TCP case (right set of curves), the period is commensurate with the round-trip time of the TCP flow on the first hop, whereas in the other case (left curves), the periods were chosen to be simple multiples.

3.4 From Delay to Jitter

So far we have considered positive functions f which act on Z at a single time point only. In fact more general functions of the form $f(Z(0), Z(t_1), \dots, Z(t_k))$ can be considered, which gives access to the temporal structure of Z . Key examples are the n -dimensional distributions of the process, and the *delay variation* or *jitter*.

Palm Calculus can deal with this greater generality by considering clusters of (non-intrusive) probes sent at epochs $\{T_n\}$ that form a stationary and ergodic point process. Each cluster consists of $k + 1$ probes sent at times $T_n + t_i$, $i = 0, \dots, k$ with $t_0 = 0$. Palm calculus can then be applied by formulating the clusters as marks, the probe process thereby becoming a *marked* point process (for details see [1]). As before one can then measure without bias the average behaviour of any such function, that is:

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N f(Z(T_n), \dots, Z(T_n + t_k)) = E[f(Z(0), \dots, Z(t_k))].$$

For example, we show how to measure jitter on a time scale of τ , that is, we desire the distribution of $J_\tau(t) = Z(t + \tau) - Z(t)$. Let the clusters arrive as a renewal process with inter-arrivals distributed uniformly over $[9\tau, 10\tau]$. This process is mixing. Each cluster will consist of two points, the cluster seed at T_n , and a trailing probe at $T_n + \tau$. We then simply collect the jitter values $\{J_\tau(T_n)\}$ and estimate its distribution by forming a histogram (technically, this implies defining multiple functions f , each an indicator function for a histogram bin, and counting the hits in each. These counts are positive, although jitter itself takes either sign).

4. INTRUSIVE MEASUREMENT

The last section dealt with the non-intrusive case. Here we consider implications for inversion and sampling bias arising from ‘real’ probes of positive size. We then combine our findings to conclude that Poisson probing should no longer be regarded as a viable default choice, and suggest an alternative.

The key new element is that now probes influence system evolution. This does not affect the existence of the virtual delay process $W(t)$ (what a zero sized observer would see when arriving to the system), nor our final aim, namely to measure the ground truth $Z^*(t)$, the delay that would be observed by a positive sized probe arriving to the *unperturbed* system at time t . However there are several important changes:

- (i) our observable $Z(t)$ is no longer $W(t)$, as the available data is now what a real probe of positive size would experience;
- (ii) observations of $Z(t)$ are not observations of $Z^*(t)$: an inversion issue arises;
- (iii) each of $Z(t)$, $Z^*(t)$ and $W(t)$ now depends on the probe process, in particular the probe size.

Our goal is to estimate $Z^*(t)$ using observations of $Z(t)$ at the probing times T_n , and to determine in what sense, if at all, the following modified form of Equation (4) holds:

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N f(Z(T_n)) = E[f(Z^*(0))]. \quad (7)$$

Notice that this equation defaults to Equation (4) if probes are non-intrusive, in which case $Z(t) = Z^*(t) = W(t)$.

4.1 PASTA

We revisit PASTA with an emphasis on probing applications. The general setting of Section 3.1 continues to hold, with the changes as listed above.

In the intrusive case, sampling bias is typically present. For instance, consider the ‘Uniform’ renewal process with support on $[0.9\mu, 1.1\mu]$ in the left hand plot in Figure 5. The negative bias results from the probes only weakly seeing the contribution to load of other probes, which arrive at least 0.9μ from them.

There is no reason for Equation (5) to hold in general. However, the PASTA property of Wolff [22] states that, for Poisson probes, it




Figure 8: Simulation illustrating the validity of PASTA in a multi-hop system, albeit with inversion bias, for 4 different packet sizes (intrusiveness levels).

does, provided the Lack of Anticipation Assumption (LAA) holds. However, we are not told which network scenarios satisfy LAA, and therefore when PASTA will hold for real probing involves many open questions. We therefore provide the following result and sketch of proof, which states that PASTA does in fact hold for delay in the traditional path model.

THEOREM 3. *Under the above assumptions, for Poisson probes (intrusive or not), in a FIFO tandem queueing network,*

$$P_0(Z(0) \in A) = P(Z(0) \in A). \quad (8)$$

Sketch of proof: In a FIFO tandem queueing network, $Z(t)$ is predictable w.r.t. the filtration $\{\mathcal{F}_t\}$, where \mathcal{F}_t is the smallest σ -field that contains the past of the probe process before time t and the (independent) σ -field that contains the whole history of the CT marked point processes (see [1], sections 1.8, 3.3).

Further, since a Poisson process is mixing, Theorem 2 is true and so the product shift is ergodic. Using Equation (8) then allows the same steps as in Section 3.3 to carry through, yielding

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N f(Z(T_n)) = E[f(Z(0))]. \quad (9)$$

This statement of PASTA in our setting reaffirms the fact that what it provides is unbiased sampling of (functions of) the total system Z . This says nothing about Equation (7), which includes the inversion step taking us back to our target, Z^* .

Figure 8 illustrates PASTA using a 3-hop network with capacities of [2, 20, 10] Mbps. Each hop carries 1-hop-persistent cross-traffic which is of type [periodic,Pareto,TCP], a combination which includes long-range dependence, and potential for phase-locking due to either the TCP (again a single flow) or periodic streams. Again for space reasons we omit details. Delay marginals, obtained from 50000 probes, are plotted over a range of intrusiveness, achieved with 4 different probe sizes. The results show, as we expect, that PASTA continues to hold for delay, despite the dangerous periodic components of cross-traffic, but that inversion bias is a problem.

4.2 Controlling Intrusiveness

One strategy to reach the objective of Equation (7) is to use a minimally perturbative stream to reduce bias (e.g., [23]). When packet sizes cannot be made arbitrarily small, non-intrusiveness can be achieved by ‘rare probing’, i.e., making the inter-arrival times large. The aim of the present section is to give a justification of this strategy in a limited but natural setting.

Setting: We assume that the queueing system without probes is described by a Markov kernel in continuous time, H_t , on some denumerable state space \mathcal{E} that is irreducible and positive recurrent with stationary probability π . We assume that the transmission of a probe is represented by applying another Markov kernel K ; more precisely, if the state of the system just before a probe is sent is described by the probability measure ν on \mathcal{E} , then the law of the state of the system when this probe reaches the receiver is νK (also on \mathcal{E}). Probe $n + 1$ is sent a random time $a\tau$ after n is received, where $a \in \mathbb{R}^+$ is a scaling factor and τ has law I on \mathbb{R}^+ . Hence, the probes are not renewal.

The Markov kernel of the total system (describing the law of the system just before probes are sent) is given by

$$\widehat{P}_a = K \int_t H_{at} I(dt).$$

We also assume that \widehat{P}_a is irreducible and positive recurrent on \mathcal{E} , with stationary probability π_a , provided a is large enough.

Rare Probing: By definition, rare probing is obtained when a goes to infinity. Within the setting of this subsection, the LHS of Equation (7) is a.s. equal to $\sum_i \pi_a(i) f(i)$ (thanks to the ergodic theorem for Markov chains), whereas its RHS is $\sum_i \pi(i) f(i)$. Our main result (proved in the appendix) on rare probing is:

THEOREM 4. *Assume that (i) The Markov kernel H_t is s.t. the parameters of the exponential sojourn times in the states of \mathcal{E} are uniformly bounded from above. (ii) The Markov kernel J of the embedded chain of H_t is α -Doebelin for some $0 < \alpha < 1$. (3) The law I has no mass at 0 (i.e. probe transmissions are always separated). Then Equation (7) holds in the following asymptotic sense: for all bounded functions f , for all $\epsilon > 0$, there exists an A such that for all $a > A$,*

$$\left| \sum_i \pi_a(i) f(i) - \sum_i \pi(i) f(i) \right| < \epsilon. \quad (10)$$

Notice that the above limit implies that both inversion and sampling bias go to zero.

Our theorem, valid for delay only, and under a particular set of conditions, is only an indication of the kind of result that could be obtained based on ‘rarity’. There is much scope for future work in this technically challenging area.

4.3 A Replacement for Poisson Probing

In active probing practice, PASTA has been employed either as a default choice, or in the belief that it will lead to zero estimation bias, and always in the expectation that it can do no harm. We claim that there is little justification in applying PASTA in the overwhelming majority of situations, and consequently that Poisson probing as a default is both undeserved and misleading, and should be replaced. We recap the case against Poisson probing before suggesting what it could be replaced *with*. As a detailed summary already appears in the introduction, here we present an account focusing on practical implications.

Inflexibility: NIMASTA and rare probing show that there is a large class of probing streams which avoid both sampling bias (if intrusiveness is controlled) and the possibility of phase-locking. Poisson probing is a lost opportunity to explore this space to minimize inversion bias as well as variance.

Inapplicability to probe patterns: The techniques of active probing are rarely based on isolated probes. Section 3.4 showed that NIMASTA is useful for (mixing) probe patterns, which could be crafted to achieve given measurement aims. Naturally, Poisson probes cannot form patterns with desired properties. Exponentially distributed probe pattern inter-arrivals have been suggested

for packet pairs [20], and packet trains [9]. However, there is no theoretical rationale for this. In active probing PASTA applies only to a stream of Poisson *packets*, and cannot justify any inference based on temporal behaviour between probes of a pair, where interactions are not memoryless.

Beyond delay, inversion bias dominates: The techniques of active probing frequently look beyond end-to-end delay. As the target $Z^*(t)$ moves from the observable itself to quantities which are far removed from it (such as details of CT at individual hops), the inversion task becomes increasingly complex. Even in the non-intrusive case, complex inversion is required (which if poorly performed results in large inversion bias) except in very special cases such the system end-to-end delay. PASTA does not solve this problem, indeed Poisson probing prevents probe patterns from being designed which could help circumvent the hard inversion problem.

There is no single probing guideline which offers optimality for all probing purposes, and in particular, it is beyond the scope of this paper to derive optimal probe patterns. Instead, we wish to provide a new default for probe sending to replace exponential separation in contexts where the sending of many separate probe patterns applies. We seek the following features:

- Suitable for separating probe patterns, not only probes,
- Generic, to allow freedom in probe pattern design: we specify pattern separations, not the entire process,
- Compatible with mixing, to reduce the risk of phase-lock,
- Tunable, to enable bias/variance tradeoffs.

The following choice satisfies these criteria:

Probe Pattern Separation Rule: *select inter-probe (or probe pattern) separations as i.i.d. positive random variables, with a distribution which contains an interval where the density is bounded above zero, and whose support is lower bounded away from zero.*

In the special case where each probe pattern is a single probe, this separation rule fully specifies a probing arrival process, namely a mixing renewal process. An example is given by the renewal process with inter-arrivals uniformly distributed over $[0.9\mu, 1.1\mu]$ of Figure 6. Since the support is lower bounded at 0.9μ , probes are guaranteed to be well separated, enabling them to act more independently, thereby reducing variance. Whereas the mean separation μ can be used to control probe rarity, the lower bound (and distribution shape) can be tuned to tradeoff sampling bias, inversion bias and variance.

In the case of probe patterns, the rule does not by itself specify the probe process, nor, therefore, determine all of its properties. However, selecting separations to be i.i.d. inherently ‘encourages’ mixing (recall section 3.4 in the non-intrusive case). Again, the enforced separation helps the patterns to make uncorrelated measurements, reducing variance. Finally, note that the support can be tuned close to zero if necessary, so Poisson-like probing is not excluded, for the (rare) cases where this is truly appropriate.

5. RELATED WORK

The problem of identifying the conditions under which observations of a stochastic system coincide with the stationary distribution of the observed process has a long history, starting with Descloux in 1967 [4]. Wolff named, gave the first rigorous proof for, and popularized the PASTA principle [22], although that principle was known earlier. Refer to [11] for a thorough review of pre-1990 work. Melamed and Whitt [11] later derived conditions for ASTA (Arrivals that See Time Averages) to hold.

The use of PASTA to justify network measurements of at Poisson epochs was pioneered by Paxson. He [13, 14], carried out mea-

surements of the ‘routing state’, delays, losses and TCP bulk transfers at Poisson epochs. The IETF IP Performance Metrics (IPPM) Group [7] built upon PASTA and Paxson’s results to recommend the use of Poisson sampling, for example in RFC 2330 [15]. RFC 2330 also observes that non-Poisson probes such as uniform, geometric, additive random, or other probes can be used for a variety of practical reasons; for example, the interval between Poisson probes can be arbitrarily large or small, and such probes cannot be implemented in real systems, hence the need to use implementable and ‘close enough to Poisson’ probes such as truncated Poisson probes. Note that all the probing processes mentioned above have the mixing property required for the NIMASTA theorem to hold.

Since then, Poisson probing has become part of the conventional wisdom of network measurements [23, 20, 9]. Recent work has attempted to better understand the impact of PASTA and the design of estimators for active probing. Bin Tariq et al. [12] empirically examined the difference between Poisson and periodic sampling, and show that, in many cases, the difference between estimates of delay and loss obtained with Poisson and periodic probes are not significant. Sommers et al. [19] set out to understand the probing process best suited to measure packet loss. They propose the use of a geometrically distributed packet pair to estimate the duration of loss periods better than can be done with Poisson probes. Roughan analyzed the Poisson probing of queueing systems and the impact of correlations in the observed (delay) process [18, 17]. Roughan’s work is another step towards developing estimators that take not just bias but also variance into account.

There is a substantial literature on *perturbation analysis* (see [6]) which addresses the ‘reversed’ problem of determining the behavior of a perturbed system from that of the unperturbed one. However, there are no immediate or simple answers to the difficult inversion problems of active probing, for example as explored in [10].

6. CONCLUSION

We provide a rigorous analysis of the precise role and relevance of PASTA in designing estimators for network measurement. Conventional wisdom holds that Poisson sampling must be used to avoid bias, and the PASTA property is cited as justification. Using end-to-end delay as a tractable example, we show that this is simplistic, and ignores two important crucial aspects of probing - inversion to the parameter of interest, and intrusiveness. We derive ‘NIMASTA’, the fact that zero sampling bias is shared by a large set of mixing probing processes in the non-intrusive case, which also avoid phase-lock problems. The zero estimation bias property is exclusive to Poisson only when probes are intrusive. However, it is extremely challenging to invert from observed delay values to the target parameter, resulting in inversion bias, which can dominate over sampling bias. Hence, even with Poisson probes, it is desirable to make them non-intrusive, by sending them as rarely as possible, and we give a theorem making this precise. In the rare case however (as well as generally) the choice of the best probing stream is dependent on the total bias (sampling and inversion) and variance, and furthermore, we explain why PASTA cannot be used to justify exponentially separated packet pairs or trains. We conclude that Poisson probing is rarely indicated, and recommend a *Probe Pattern Separation Rule* as an alternative default for both probes and probe patterns, which offers several advantages.

7. APPENDIX

Sketch of proof of Theorem 4 The idea of the proof is: if probes are separated by a time which is long compared to the ‘speed of convergence to steady state’ of the unperturbed system, then probes

should sample a distribution close to that of the unperturbed system. To make this rigorous, we need assumptions on the speed of this convergence. The Doeblin property is one such assumption.

Let $\alpha \in (0, 1)$. A Markov P kernel is α -Doeblin if it can be represented in the form $P = (1 - \alpha)A + \alpha Q$ where Q is some Markov kernel on \mathcal{E} and A is a rank 1 Markov kernel, namely the matrix A has all its lines equal. A sufficient condition for a Markov kernel P to be α -Doeblin is that $\forall j, P(i, j) \geq (1 - \alpha)\phi(j)$ uniformly on i , for some probability measure ϕ on \mathcal{E} . Our assumption is that the Markov kernel J of the embedded chain of H_t is s.t. J^n is α -Doeblin for some constant $0 < \alpha < 1$ and some integer $n > 0$.

The proof relies on the following classical properties, where $\|\cdot\|$ denotes the L_1 norm. (1) All Markov kernels P on \mathcal{E} are non-expansive for the L_1 norm in the sense that if ν and ν' are probability measures, then $\|\nu P - \nu' P\| \leq \|\nu - \nu'\|$. (2) All α -Doeblin Markov kernels P are α -contracting for that norm, that is such that for all ν and ν' as above $\|\nu P - \nu' P\| \leq \alpha\|\nu - \nu'\|$ which is immediate from the representation $P = (1 - \alpha)A + \alpha Q$ and (1). (3) All α -Doeblin Markov kernels P that admit an invariant probability κ are such that for all probability measures ν , $\|\nu P^n - \kappa\| \leq \alpha^n\|\nu - \kappa\|$ (this follows from (2)). This uniform contraction property gives the uniform speed of convergence alluded to above. (4) If K is an arbitrary Markov kernel on \mathcal{E} and H is α -Doeblin, then KH and HK are both α -Doeblin.

We will also use the following lemma, which shows that any probability measure which is "nearly invariant" w.r.t. an α -Doeblin Markov kernel is close to the invariant measure of this kernel:

LEMMA 7.1. *If P is an α -Doeblin Markov kernel with stationary distribution π and if ν is a probability measure such that $\|\nu - \nu P\| \leq \epsilon$, then $\|\pi - \nu\| \leq \epsilon/(1 - \alpha)$.*

The proof follows from: $\|\pi - \nu\| = \|\pi - \nu P + \nu P - \nu\| \leq \|\pi P - \nu P\| + \|\nu P - \nu\| \leq \alpha\|\pi - \nu\| + \epsilon$.

The first step to prove Theorem 4 is to show that, $\forall a$ large enough, $\hat{H}_a = \int H_{at} I(dt)$ is β -Doeblin for some coefficient β that does not depend on a . Let B be a Borel set of \mathbb{R} such that $I(B) > 0$ and such that b , the infimum of the points of B , is positive (such a set exists because of our assumption on I). In view of the exponential law of the sojourn times in the state of the continuous time Markov chain $\{H_t\}$, there is a probability $q > 0$ that there are n jumps of the chain in the interval $[0, b]$. Hence for $a > 1$:

$$\begin{aligned} \hat{H}_a &= \int H_{at} I(dt) \geq \int_B H_{at} I(dt) \\ &= \int_B H_b H_{at-b} I(dt) = H_b \int_B H_{at-b} I(dt) \geq q J^n G_a, \end{aligned}$$

where G_a is the sub-Markov kernel $\int_B H_{at-b} I(dt)$. Since we assumed J^n to be α -Doeblin we get that, $\hat{H}_a \geq q J^n G_a \geq q(1 - \alpha)AG_a = \beta B_a$, where $B_a = AG_a$ is a substochastic matrix of rank 1 and where $0 < \beta = q(1 - \alpha) < 1$ does not depend on a , which concludes the proof of the first step in view of Property 4.

The second step consists in showing that for all probability measures ν on \mathcal{E} , $\lim_{a \rightarrow \infty} \|\pi - \nu \hat{H}_a\| = 0$. Since $\{H_t\}$ is Doeblin,

$$\lim_{a \rightarrow \infty} \|\nu H_{at} - \pi\| = 0 \quad \forall t > 0. \quad (11)$$

This concludes the proof of the second step since we have

$$\begin{aligned} \lim_{a \rightarrow \infty} \|\pi - \nu \hat{H}_a\| &= \lim_{a \rightarrow \infty} \left\| \int_t (\pi - \nu H_{at}) I(dt) \right\| \leq \\ \lim_{a \rightarrow \infty} \int_t \|\pi - \nu H_{at}\| I(dt) &= \int_t \lim_{a \rightarrow \infty} \|\pi - \nu H_{at}\| I(dt) = 0, \end{aligned}$$

where we used Lebesgue's dominated convergence theorem to interchange the integral and the limit, and then (11).

The first step and Property 4 imply that $P_a = K \hat{H}_a$ is β -Doeblin. When choosing $\nu = \pi K$ in the second step, we get that $\forall \epsilon > 0$, $\|\pi - \nu \hat{H}_a\| = \|\pi - \pi P_a\| \leq \epsilon$ for a large enough. Hence, Lemma 7.1 implies that for a large enough, $\|\pi - \pi_a\| \leq \frac{\epsilon}{1 - \beta}$. The result then follows from immediate bounds. \square

Enhancements to NS-2 to measure ground truth $Z^*(t)$

Using the traces of all arrivals and departures from a single hop, we store the queue size $W_h(t)$ of hop h at any time t by exploiting the fact that it is piecewise linear. The $W_i(t)$ are combined over hops to calculate $Z_p(t)$, the delay that a packet of size p injected at an arbitrary time t would have experienced. If D_h denotes the propagation delay of hop h , then $Z_p(t)$ is given recursively by

$$\begin{aligned} Z_p(t) &= W_1(t) + p/C_1 + D_1 + \\ &W_2(t + W_1(t) + p/C_1 + D_1) + p/C_2 + D_2 + \\ &W_3(\dots) \dots \quad \text{to the last hop.} \end{aligned} \quad (12)$$

8. REFERENCES

- [1] F. Baccelli and P. Bremaud. *Elements of Queueing Theory*. Springer Verlag, Applications of Mathematics, Second edition, 2003.
- [2] D. R. Cox. in *Long-range dependence: a review*, H.A. David and H.T. David editors, chapter Statistics : an Appraisal, pages 55–74. Iowa State University Press, Ames (IA), 1984.
- [3] D. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes*. Springer-Verlag, 1988.
- [4] A. Descloux. On the Validity of a Particular Subscriber's View. In *Proc. of Fifth International Teletraffic Congress*, 1967.
- [5] D. Gaver and P. Lewis. First-order Autoregressive Gamma Sequences and Point Processes. *Adv. Appl. Prob.*, 12:727–745, 1980.
- [6] P. Glasserman. *Gradient Estimation via Perturbation Analysis*. Kluwer Academic Publishers, Boston, 1991.
- [7] IETF IP Performance Metrics (IPPM). <http://www.ietf.org/html.charters/ippm-charter.html>, 2005.
- [8] L. Kleinrock. *Queueing Systems*, Volume I: Theory. John Wiley and Sons, 1975.
- [9] X. Liu, K. Ravindran, B. Liu, and D. Loguinov. Single-Hop Probing Asymptotics in Available Bandwidth Estimation: Sample-Path Analysis. In *Proc. of IMC*, October 2004.
- [10] S. Machiraju, D. Veitch, F. Baccelli, and J. Bolot. Theory and Practice of Cross-traffic Estimation via Probes. Technical Report RR-5763, INRIA-ENS, 2005.
- [11] B. Melamed and W. Whitt. On arrivals that see time averages. *Operations Research*, 38(1):156–172, 1990.
- [12] M. Tariq, A. Dhamdhere, C. Dovrolis, and M. Ammar. Poisson versus Periodic Path Probing (or, Does PASTA Matter)? *Proc. IMC*, 2005.
- [13] V. Paxson. End-to-end routing behavior in the internet. In *Proc. of ACM SIGCOMM*, 1997.
- [14] V. Paxson. End-to-end Internet packet dynamics. *IEEE/ACM Transactions on Networking*, 7(3):277–292, 1999.
- [15] V. Paxson, G. Almes, J. Madhavi, and M. Mathis. RFC 2330 - Framework for IP Performance Metrics, May 1998.
- [16] K. Petersen. *Ergodic Theory*. Cambridge University Press, Cambridge England, 1983.
- [17] M. Roughan. A Comparison of Poisson and Uniform Sampling for Active Measurements. *To appear in IEEE JSAC*.
- [18] M. Roughan. Fundamental Bounds on the Accuracy of Network Measurements. In *Proc. of ACM SIGMETRICS*, June 2005.
- [19] J. Sommers, P. Barford, N. Duffield, and A. Ron. Improving Accuracy in End-to-End Loss Measurement. In *Proc. of ACM SIGCOMM*, 2005.
- [20] J. Strauss, D. Katabi, and F. Kaashoek. A Measurement Study of Available Bandwidth Estimation Tools. In *Proc. of IMC*, 2003.
- [21] NS-2(Network Simulator) <http://www.isi.edu/nsnam/ns/>.
- [22] R. Wolff. Poisson Arrivals see Time Averages. *Operations Research*, 30(2):223–231, 1982.
- [23] Y. Zhang, N. Duffield, and V. Paxson. On the Constancy of Internet Path Properties. In *Proc. of IMW*, 2001.