REPORTRAPPORT

PNA

Probability, Networks and Algorithms

Probability, Networks and Algorithms

Resource dimensioning through buffer sampling

M.R.H. Mandjes, R. van de Meent

Centrum voor Wiskunde en Informatica (CWI) is the national research institute for Mathematics and Computer Science. It is sponsored by the Netherlands Organisation for Scientific Research (NWO).
CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics.

CWI's research has a theme-oriented structure and is grouped into four clusters. Listed below are the names of the clusters and in parentheses their acronyms.

Probability, Networks and Algorithms (PNA)

Software Engineering (SEN)

Modelling, Analysis and Simulation (MAS)

Information Systems (INS)

Resource dimensioning through buffer sampling

ABSTRACT

Link dimensioning, i.e., selecting a (minimal) link capacity such that the users' performance requirements are met, is a crucial component of network design. It requires insight into the interrelationship between the traffic offered (in terms of the mean offered load M, but also its fluctuation around the mean, i.e., 'burstiness'), the envisioned performance level, and the capacity needed. We first derive, for different performance criteria, theoretical dimensioning formulae that estimate the required capacity C as a function of the input traffic and the performance target. For the special case of Gaussian input traffic these formulae reduce to $C = M + \text{alpha } V$, where alpha directly relates to the performance requirement (as agreed upon in a service level agreement) and V reflects the burstiness (at the timescale of interest). We also observe that Gaussianity applies for virtually all realistic scenarios; notably, already for a relatively low aggregation level the Gaussianity assumption is justified. As estimating M is relatively straightforward, the remaining open issue concerns the estimation of V. We argue that, particularly if V corresponds to small time-scales, it may be inaccurate to estimate it directly from the traffic traces. Therefore, we propose an indirect method that samples the buffer content, estimates the buffer content distribution, and 'inverts' this to the variance. We validate the inversion through extensive numerical experiments (using a sizeable collection of traffic traces from various representative locations); the resulting estimate of V is then inserted in the dimensioning formula. These experiments show that both the inversion and the dimensioning formula are remarkably accurate.

Resource dimensioning through buffer sampling

Michel Mandjes and Remco van de Meent

Abstract

Link dimensioning, i.e., selecting a (minimal) link capacity such that the users' performance requirements are met, is a crucial component of network design. It requires insight into the interrelationship between the traffic offered (in terms of the mean offered load M, but also its fluctuation around the mean, i.e., 'burstiness'), the envisioned performance level, and the capacity needed. We first derive, for different performance criteria, theoretical dimensioning formulae that estimate the required capacity C as a function of the input traffic and the performance target. For the special case of Gaussian input traffic these formulae reduce to C $= M + \alpha V$, where α directly relates to the performance requirement (as agreed upon in a service level agreement) and V reflects the burstiness (at the timescale of interest). We also observe that Gaussianity applies for virtually all realistic scenarios; notably, already for a relatively low aggregation level the Gaussianity assumption is justified.

As estimating M is relatively straightforward, the remaining open issue concerns the estimation of V. We argue that, particularly if V corresponds to small time-scales, it may be inaccurate to estimate it directly from the traffic traces. Therefore, we propose an indirect method that samples the buffer content, estimates the buffer content distribution, and 'inverts' this to the variance. We validate the inversion through extensive numerical experiments (using a sizeable collection of traffic traces from various representative locations); the resulting estimate of V is then inserted in the dimensioning formula. These experiments show that both the inversion and the dimensioning formula are remarkably accurate.

Index Terms

Network dimensioning ● Quality-of-Service ● buffer sampling ● Gaussian traffic ● inversion ● large deviations

I. Introduction

ADEQUATE resource dimensioning requires a thorough insight into the interrelationship between: (i) the traffic offered (in terms of the average load, but also its fluctuations), (ii) the desired level of performance, (iii) the required capacity. It is clear that more capacity is needed when the offered load becomes higher, the fluctuations become fiercer, or the performance criterion becomes more stringent. However, to make precise predictions about the *amount* of capacity that should be added, advanced modeling and performance techniques are required. These predictions are of crucial importance, as scarce dimensioning inevitably leads to performance degradation, whereas 'generous' dimensioning policies essentially result in a waste of resources.

Michel Mandjes (`mmandjes@science.uva.nl`) is with the University of Amsterdam, Plantage Muidergracht 24, 1018 TV, Amsterdam, the Netherlands, and CWI, P.O. Box 94079, 1090 GB Amsterdam, the Netherlands; Remco van de Meent (`r.vandemeent@utwente.nl`) is with the University of Twente, P.O. Box 217, 7500 AE, Enschede, the Netherlands.

In the context of dimensioning IP nodes, several additional considerations play a role. IP networks usually carry a broad variety of traffic types, of which some can tolerate relatively substantial latency (email, particular web applications), whereas others have strict delay requirements (interactive services). As argued by e.g. Fraleigh *et al.* [12], two approaches can be followed: (i) actively discriminating by traffic differentiation mechanisms (i.e., preferential treatment for the more demanding traffic types), or (ii) provisioning sufficiently resources so that all traffic types meet their performance requirements (i.e., without any traffic differentiation). In the former option, the cost associated to the management and operation of the network are relatively high, while the efficiency gain (compared to the latter option) is usually modest (particularly when the level of aggregation is high). Therefore, following [12], the focus in this paper is on dimensioning as the approach for delivering performance requirements. (As an aside, we mention that even if one opts for applying traffic-differentiation, there is often still a need for dimensioning as well: for instance, as argued in [5], if *diffserv* is applied, then still bandwidth dimensioning for traffic aggregates needs to be done. Put differently, the distinction between traffic differentiation and resource dimensioning is not sharp).

Above we identified three crucial prerequisites for link dimensioning: (i) a reliable traffic model, (ii) a performance target to be met, and (iii) formulae computing the performance for a given traffic model and given network resources. Having these at our disposal, we can find expressions for the minimum link capacity required in order to offer a traffic stream with given characteristics a given performance target; we refer to these as *dimensioning formulae*. In *Section II* we present a number of generic dimensioning formulae; generic in the sense that they are valid for any (stationary) traffic stream. These formulae, however, are still quite implicit; they require knowledge of the full moment generating function (mgf) of the traffic offered at any time-scale, something that is typically hard to measure and estimate. To overcome this problem it would be helpful if we could restrict ourselves to specific classes of, still sufficiently general and versatile, random processes.

Two such classes of traffic models have been discussed extensively in the literature. The first is what could be called 'flow-oriented', where *flows* could be TCP connections, UDP streams, etc. In this approach flows are modeled (arrival rate, duration, traffic rate while being active, etc.); it is noted that tools like NetFlow (in Cisco routers) already allow for measurement of some flow characteristics (e.g., size and duration). Several papers have found, in different situations, accurate flow-oriented traffic models, see for instance [2], [3]. An alternative is to model the *aggregate* stream, rather than individual flows. *Gaussian models* (of which fractional Brownian motion is a special case) form a class of models that is specifically suitable for describing highly aggregated streams. The Gaussian model consists of (i) a *mean rate M* (to which we also refer as 'load') and (ii) a *variance curve $V(\cdot)$* (so that $V(t)$ corresponds to the variance of the amount of traffic offered in an arbitrary time window of length $t \geq 0$). Having observed that the dimensioning formulae of Section II greatly simplify for Gaussian traffic, in *Section III* we further assess in detail to what extent one can assume Gaussianity. We discuss which factors affect the Gaussianity (the length of the measurement interval, the level of aggregation), and conclude that in virtually any representative situation (where we considered various locations, and various points in time) the Gaussian model fits well. We emphasize that, particularly under high load, flow-oriented models and aggregate-stream models do not exclude one another; we come back to this issue later.

Assuming Gaussianity, the dimensioning formulae of Section II require estimates of the mean traffic rate M and the variance curve $V(\cdot)$ to find the required bandwidth for a given performance target. Estimating M is relatively straightforward, and can be done through rough traffic measurements (for instance over 5-minute intervals). Estimating the variance curve, however, could be substantially harder: particularly on smaller time-scales, it is hard to do accurate measurements through SNMP (simple network management protocol). *Section IV* presents a novel, efficient technique for estimating $V(\cdot)$, by coarse-grained sampling the buffer content, estimating the buffer content distribution, and 'inverting' this into the variance curve. Importantly, this procedure eliminates the need for traffic measurements on small time-scales; instead we measure (for instance at a constant frequency, but this is by no means necessary) the buffer content. In this sense, we remark that the procedure we propose is rather counter-intuitive: one would expect that one needs measurements of the traffic offered in intervals of length T to accurately estimate $V(T)$, but apparently one can alternatively sample the buffer content. In fact, one of the attractive features of our 'inversion approach' is that it yields the *entire* variance curve $V(\cdot)$ (evidently up to some finite time horizon), rather than just $V(T)$ for some pre-specified T.

Section V assesses the accuracy of the inversion approach, through simulation experiments with both synthetic traffic and real network traces. These validations show excellent performance, in the sense that the $V(\cdot)$-curve is indeed estimated well, even if the traffic stream deviates considerably from Gaussian. We also investigate the required measurement effort (in terms of number of samples, and the time between two subsequent samples) in order to obtain a reliable estimate of the buffer content distribution.

The next step is to insert the estimated variance into the dimensioning formulae of Section II. In *Section VI* we do so for our reference set of real traces. We compare the resulting capacity by its 'empirical counterpart', i.e., the minimum link rate such that for the given trace the performance requirement is met. We also systematically study the impact of the performance criterion on the required link rate, and provide implementation guidelines.

Section VII concludes the paper. We discuss the results, and the applicability of our approach. We also reflect on a number of related methods.

II. GENERIC DIMENSIONING FORMULAE

As explained in the introduction, an important prerequisite for dimensioning are formulae that determine, for given characteristics of the offered traffic and performance target, the minimum required link rate. Preferably, these dimensioning formulae have minimal requirements on the 'nature' of the traffic offered; for instance, we do not want to impose any conditions on its correlation structure.

In this section, we present formulae that we derive under extremely weak conditions on the traffic process. The only substantial assumption is that we require that the traffic stream be stationary: with $A(s,t)$ denoting the amount of traffic arrived in $[s,t)$, it is assumed that the distribution of $A(s+\delta,t+\delta)$ does not depend on δ (but just on the interval length $t-s$). In the sequel we use the abbreviation $A(t) := A(0,t)$. In this paper we study dimensioning with respect to two performance criteria: (A) 'link transparency' and (B) buffer overflow.

(A) In 'link transparency', cf. [5], the main objective of bandwidth dimensioning is to ensure

that the links are more or less 'transparent' to the users, in that the users should not (or almost never) perceive any performance degradation due to a lack of bandwidth. Clearly, this objective will be achieved when the link rate is chosen such that only during a small fraction of time ε the aggregate rate of the offered traffic (measured on a sufficiently small time scale T) exceeds the link rate: $\mathbb{P}(A(T) \geq \text{C}T) \leq \varepsilon$. The values to be chosen for the parameters T and ε typically depend on the specific needs of the application(s) involved. Clearly, the more interactive the application, the smaller T and ε should be chosen; network operators should choose them in line with the service level agreements they agreed on with their clients.

Given the criterion $\mathbb{P}(A(T) \geq \text{C}T) \leq \varepsilon$, we now derive a formula for the minimal link rate needed. Relying on the celebrated Chernoff bound, we have

$$\mathbb{P}(A(T) \geq \text{C}T) \leq \inf_{\vartheta \geq 0} \left(e^{-\vartheta \text{C}T} \mathbb{E} e^{\vartheta A(T)} \right), \text{for any } \vartheta \geq 0.$$

As this inequality holds for *any* $\vartheta \geq 0$, we see that, in order to be sure that $\mathbb{P}(A(T) \geq \text{C}T) \leq \varepsilon$ it suffices to take the link rate C larger than

$$\text{C}_1 := \inf_{\vartheta \geq 0} \frac{\log \mathbb{E} \exp(\vartheta A(T)) - \log \varepsilon}{\vartheta T}. \tag{1}$$

(B) The link-transparency criterion did not explicitly take into account the option of buffering packets. The distribution of the steady-state buffer content Q can be expressed in terms of the arrival process $A(\cdot)$: Reich's formula says that Q is distributed as the maximum of the process $A(-t, 0) - \text{C}t$:

$$Q =_{\text{d}} \sup_{t \geq 0}(A(-t, 0) - \text{C}t),$$

where '$=_{\text{d}}$' denotes equality in distribution. A second performance criterion could be to choose C such that $\mathbb{P}(Q > \text{B}) < \varepsilon$, where B is the router's buffer size (alternatively, B/C can be interpreted as an upper bound on the delay); observe that Q depends on C.

In order to find the minimum required capacity, we need to characterize $\mathbb{P}(Q > \text{B})$ as a function of C. This is done as follows. In the first place, observe that 'Reich' entails that $\mathbb{P}(Q > \text{B})$ is the probability that $A(-t, 0)$ exceeds $\text{B} + \text{C}t$ for some $t \geq 0$. This is a union of events (namely, the union over $t \geq 0$ of the events $\{A(-t, 0) \geq \text{B} + \text{C}t\}$); it often is accurate to approximate the probability of a union of events by the largest of the individual probabilities. We thus obtain

$$\mathbb{P}(Q > \text{B}) \equiv \mathbb{P}(\exists t : A(-t, 0) \geq \text{B} + \text{C}t)$$
$$\approx \sup_{t \geq 0} \mathbb{P}(A(-t, 0) \geq \text{B} + \text{C}t),$$

see, for instance, [12]; a further large-deviations justification for the 'principle of the largest term' is given in e.g. [1], [6]. It can be argued, see for instance Section 10.3 of [13], that the Chernoff bound is in fact a reasonable approximation, so that we obtain

$$\mathbb{P}(Q > \text{B}) \approx \sup_{t \geq 0} \inf_{\vartheta \geq 0} \left(e^{-\vartheta(\text{B}+\text{C}t)} \mathbb{E} e^{\vartheta A(-t, 0)} \right); \tag{2}$$

note that this approximation contains a conservative element (Chernoff bound) as well as an 'aggressive' element (principle of the largest term), and it is not clear upfront which effect dominates.

After some rearranging, we conclude that c should be at least

$$c_2 := \sup_{t \geq 0} \inf_{\vartheta \geq 0} \frac{\log \mathbb{E} \exp(\vartheta A(-t, 0)) - \log \varepsilon - \vartheta B}{\vartheta t}, \tag{3}$$

cf. for instance [16] and [27, Eq. (5)].

Clearly, application of (1) and (3) requires the estimation of the *moment generating function* (mgf)

$$\mathbb{E} e^{\vartheta A(s,t)}, \quad \vartheta \geq 0;$$

in (1) just for $s = 0$ and $t = T$, but in (3) we even need $\mathbb{E} e^{\vartheta A(-t, 0)}$ for all $t \geq 0$. Such an estimation is extremely demanding, and far from straightforward (for specific situations, the results in [11] are helpful here; we comment on these in detail in Section VI). However, imposing some additional structure on $A(\cdot)$ may greatly simplify the dimensioning formulae; this structure should of course be flexible enough to still cover all relevant traffic patterns. The following example provides such a framework: the class of Gaussian processes.

Example: Suppose that $A(\cdot)$ is a Gaussian process with stationary increments, i.e., $A(s, t)$ is normally distributed, with mean $M \cdot (t - s)$ and variance $V(t - s)$, for some mean rate $M \in \mathbb{R}$ and variance curve $V(\cdot) : \mathbb{R}^+ \to \mathbb{R}^+$. The covariance structure is fully defined by the variance function, as it holds that $\mathbb{C}\text{ov}(A(s), A(t)) = \frac{1}{2}(V(t) + V(s) - V(t - s))$. An important special case is *fractional Brownian motion*, in which $V(t)$ is proportional to t^{2H}, where $H \in (0, 1)$ is the so-called *Hurst parameter*; when choosing $H = \frac{1}{2}$, we obtain 'ordinary' Brownian motion.

When assuming that traffic is Gaussian, with $\delta := \sqrt{-2 \log \varepsilon}$, the dimensioning formulae (1) and (3) reduce to

$$c_1 = M + \frac{\delta}{T} \cdot \sqrt{V(T)} \quad \text{and} \tag{4}$$

$$c_2 = M + \inf_{t \geq 0} \left(\frac{\delta}{t} \cdot \sqrt{V(t)} - \frac{B}{t} \right), \tag{5}$$

respectively; here it is used that $A(-t, 0)$ is distributed as $A(0, t)$ (as any Gaussian process is time-reversible), and $\mathbb{E} \exp(\vartheta A(t)) = M \vartheta t + \vartheta^2 V(t)/2$ (the latter identity follows, after some calculus, from $\int e^{-\frac{1}{2} x^2} \mathrm{d}x = \sqrt{2\pi}$, and completing the square). The important consequence of this, is that for the application of the dimensioning formulae (1) and (3) it is not required anymore to estimate mgfs; instead they can be computed when we have estimates for the mean rate M and the variance curve $V(\cdot)$ at our disposal — in (1) we even need just $V(T)$, rather than the whole curve. $\diamondsuit$

III. MODELING TRAFFIC AGGREGATES BY GAUSSIAN PROCESSES

The example in the previous section showed that for Gaussian inputs the dimensioning formulae become substantially more manageable. In this section we discuss traffic models in a general context, motivate why we concentrate on Gaussian models, and assess when Gaussian models are applicable.

A. Traffic models

As argued earlier, a crucial prerequisite for using the dimensioning formulae (1) and (3) is the availability of simple, yet accurate traffic models. Clearly, the simplest model with Poisson

arrivals of packets has the undesirable feature that it fails to incorporate the (positive) correlations between packet arrivals as observed in real traces. For this reason, the model with (a superposition of) ON/OFF sources is an attractive alternative: a broad variety of correlation structures can be modeled by choosing appropriate distributions for the ON- and OFF-times. A variant of the latter model is the so-called M/G/∞ input model, in which *flows* (groups of packets, with some general distribution) arrive according to a Poisson process, remain in the system for some random time, and generate traffic during this sojourn time according to some (random or deterministic) pattern. By choosing a heavy-tailed flow-size distribution, strong positive correlations can be obtained. Two recent papers that fit these flow-level models are [2], [3].

Rather than attempting to describe traffic at the flow level, one could also opt for trying to find models for the *aggregate traffic stream*. The development of this type of models was triggered by a number of measurement studies performed in the early 1990s, such as the famous Bellcore measurements, see for instance [18], [26]. It was shown that, in many situations, the aggregate stream has self-similar properties, and is long-range dependent (i.e., has a slowly decaying autocorrelation function).

A stochastic model, advocated by Norros in [24], [25], that has many desirable properties (e.g. long-range dependency) is fBm. fBm is a self-similar process (i.e., $A(ft)$ has the same distribution as $\phi(f) A(t)$, for $f \geq 0$, and a function $\phi(\cdot)$ that does not depend on t), and, as mentioned in the previous section, falls in the class of Gaussian models. It found in recent years wide-spread use as a reference model for IP traffic; importantly $H \in (\frac{1}{2}, 1)$ corresponds to long-range dependent, positively correlated traffic. By taking $V(t) = t^{2H}$, the dimensioning formulae (4) and (5) further simplify to

$$C_1 = M + \frac{\delta}{T^{1-H}};$$

$$C_2 = M + \delta^{1/H} \left(\frac{1-H}{B} \right)^{1/H-1} H.$$

Besides the above motivations for Gaussian traffic models, their applicability can also be explained from the Central Limit Theorem (CLT). The CLT entails that the sum of a large number of 'small' independent (or weakly dependent), statistically more or less identical, random variables (users) has an approximately normal (i.e., Gaussian) distribution. Thus, one can expect that an aggregated traffic stream consisting of many individual communications may be modeled by a Gaussian stochastic process. However, it is clear that the CLT argumentation does not apply to any timescale: on the timescale of transmission of (minimum size) packets, the traffic stream is always ON/OFF (either there is transmission at link speed, or silence) — which is obviously not Gaussian. Thus, besides requiring that the number of users (referred to as 'vertical aggregation') is sufficiently high, there should also be sufficient aggregation in time ('horizontal aggregation'). Kilpi and Norros in [17] and Fraleigh *et al.* [12] pointed out the necessity of enough aggregation in both directions for traffic to be Gaussian.

We emphasize that M/G/∞ input models and Gaussian models do not exclude one another: as long as the aggregation level is sufficiently high, both models could fit very well. The intuitive reason for this is that, under those circumstances, a Poisson random variable can be approximated accurately by a Normal random variable. In light of this, it is not surprising that one can formally

Fig. 1. Distribution of the linear correlation coefficient over all measurements

prove that, in a particular limiting regime (more precisely: by speeding up the arrivals) the $\mathrm{M/G/\infty}$ input model converges to a Gaussian process; this can be proven as in [9].

We remark, however, that for the $\mathrm{M/G/\infty}$ input model it may take a substantial effort to fit all the parameters corresponding to the flow duration, traffic transmission rate, etc., see [3].

B. Gaussianity for different levels of horizontal and vertical aggregation

We now report on our findings regarding Gaussianity. These are largely in line with the conclusions in e.g. [12], [17]; more experiments can be found in [22], [23]. We have tried to make the data sets as representative as possible: the locations cover all sorts of different types of users. As we largely follow the methodology of [12], [17], we have attempted to keep this subsection brief.

The goal of this subsection is to provide empirical support for the claim that the Gaussianity assumption is justified at many locations, with very distinct types of users. The five locations we have considered are (U) a university residential network (15 traces, 1800 hosts), (R) a research institute (185 traces, 250 hosts), (C) a college network (302 traces, 1500 hosts), (A) an ADSL access network (50 traces, 2000 hosts), and (S) a server hosting provider (201 traces, 100 hosts). Each trace relates to 15 minutes (real time). At these locations, traffic is generated at average (aggregate) rates of 170, 6, 35, 120, and 12 Mbit/s, respectively.

Both [12] and [17] observe that there should be a sufficient level of aggregation to make traffic Gaussian; [12] further quantifies this claim by reporting that at least a traffic rate of 50 Mbit/sec is needed (when considering backbone links). We now verify whether this applies for our data set $\{U, R, C, A, S\}$. As a measure of Gaussianity we use the notion of the *linear correlation coefficient* as the goodness-of-fit measure, as in [15], [17]; in [21] it is shown that the use of this statistic leads to similar conclusions as the Kolmogorov-Smirnov test that was used in [12]. Fig. 1, which focuses on the timescale of 1 sec., confirms the findings of [12]: roughly spoken, for this timescale an average traffic rate of some tens of Mbit/sec seems enough to safely assume Gaussianity.

Suppose that we observe that a specific traffic stream is fairly Gaussian at a certain timescale.

Fig. 2.   Linear correlation coefficient over 5 measurements from location R

One may then wonder what this says about Gaussianity at other timescales. If Gaussianity would be (more or less) preserved across timescales, then one needs to verify Gaussianity just on one timescale. It is clear that this reasoning has its limitations; recall that at very small timescales, traffic is certainly *not* Gaussian, and also for very large timescales Gaussianity will be lost.

First, we look at an example with only a few traces. We determine the linear relation coefficient γ at 9 timescales τ_i, ranging from 5 msec. to 5 sec. The results, based on five traces from measurement location R, are plotted in Fig. 2. As reflected by the more or less horizontal lines, the graph suggests that the Gaussianity is rather constant across timescales.

Next, we investigate this for all traces. We introduce ν_γ as measure of the 'variation of the linear correlation coefficient', defined as the square root of the sample variance of the γ_{τ_i}:

$$\nu_\gamma := \sqrt{\widehat{\mathbb{V}\mathrm{ar}}\big(\gamma_{\tau_1},\ \gamma_{\tau_2},\ \ldots,\ \gamma_{\tau_9}\big)}.$$

The interpretation is that when ν_γ is low, the traffic is (more or less) equally Gaussian (or non-Gaussian) across multiple timescales.

We have computed ν_γ using all traces from each measurement location. After ordering them from low to high values of ν_γ, they are plotted in Fig. 3. Clearly, ν_γ is small in most cases: in over 95% of the traces, ν_γ is below 0.05. Thus we may conclude that γ is quite constant over different timescales; in other words: traffic that exhibits Gaussian characteristics at one timescale, is likely to be Gaussian at other timescales as well (for the timescales that we investigated, and bearing in mind the limitations mentioned above). In addition, in line with the conclusions of [12], we see from our experiments that Gaussianity usually applies from the msec-level on (in [12] even smaller timescales were considered, but these appear to be irrelevant for dimensioning purposes).

It is important to notice that in some cases, our experiments show that the Gaussian model is *not* accurate — for instance a substantial part of the traces from location R, where there is a relatively low aggregation level. As our dimensioning approach relies on the Gaussianity assumption, it is important to verify whether it still gives reasonable outcomes for situations in which the traffic is *not* Gaussian. We come back to this issue in detail in Section V.

Fig. 3. Variance of linear correlation coefficient over time, over all measurements at various timescales, at locations $\{U, R, C, A, S\}$.

IV. ESTIMATION OF THE MEAN TRAFFIC RATE AND VARIANCE CURVE

In Section II we derived the dimensioning formulae (1) and (3), which require knowledge of the mean traffic rate M and the variance curve $V(\cdot)$. As argued earlier, M can be determined by standard coarse-grained traffic measurements (e.g., polling Interfaces Group MIB counters via SNMP every 5 minutes). It is clear that determining the variance curve $V(\cdot)$ is more involved. The standard way to estimate $V(T)$ (for some given interval length T) is what we refer to as the 'direct approach': perform traffic measurements for disjoint intervals of length T, and just compute their sample variance. It is noted that the convergence of this estimator could be prohibitively slow when traffic is long-range dependent [4, Ch. I], but the approach has two other significant drawbacks:

- When measuring traffic using time windows of duration T, it is clearly possible to estimate $V(T), V(2T), V(3T)$, etc. However, these measurements obviously do not give any information on $V(\cdot)$ on time-scales *smaller* than T. Hence, to estimate $V(T)$ measurements should be done at granularity T or less. This evidently leads to a substantial measurement effort.
- The dimensioning formula (3) requires knowledge of the *entire* variance function $V(\cdot)$, whereas the direct approach described above just yields an estimate of $V(T)$ on a pre-specified timescale T. Therefore, a method that estimates the entire curve $V(\cdot)$ is preferred.

This section presents a powerful alternative to the direct approach; we refer to it as the *inversion approach*, as it 'inverts' the buffer content distribution to the variance curve. This inversion approach overcomes the problems identified above. We rely on the large-deviations framework of Section II, for Gaussian inputs, as justified in Section III.

A. Inversion formula

In this subsection, we first show how, for given variance curve $V(\cdot)$ (and mean rate M) the probability $\mathbb{P}(Q > B)$ can be approximated. Then this explicit formula is used to 'invert' this relation: we establish a formula for $V(t)$, given the (complementary) buffer content distribution $\mathbb{P}(Q > \cdot)$.

- In Section II we found the approximation (2). For the special case of Gaussian traffic, using that $\mathbb{E}\exp(\vartheta A(t)) = M\vartheta t + \vartheta^2 V(t)/2$, the minimization over ϑ can be explicitly evaluated. We find that (2) reduces to

$$\mathbb{P}(Q > \textsc{b}) \approx e^{-I(\textsc{b})}, \tag{6}$$

$$\text{with } I(\textsc{b}) := \inf_{t > 0} \frac{((\textsc{c} - M)t + \textsc{b})^2}{2V(t)}.$$

- Supposing that approximation (6) is exact, it implies that

$$\forall \textsc{b} \geq 0 : \forall t \geq 0 : -\log \mathbb{P}(Q > \textsc{b}) \leq \frac{((\textsc{c} - M)t + \textsc{b})^2}{2V(t)},$$

or, equivalently,

$$\forall t \geq 0 : \forall \textsc{b} \geq 0 : V(t) \leq \frac{((\textsc{c} - M)t + \textsc{b})^2}{-2\log \mathbb{P}(Q > \textsc{b})}.$$

Clearly, the latter inequality implies that for all $t \geq 0$,

$$V(t) \leq \inf_{\textsc{b} \geq 0} \frac{((\textsc{c} - M)t + \textsc{b})^2}{-2\log \mathbb{P}(Q > \textsc{b})}.$$

In fact, this upper bound on the variance is under mild conditions tight in a large-deviations sense (i.e., in the many-sources framework); for details we refer to Theorem 4 in [19]. Remarkably, this says that, loosely speaking, for Gaussian traffic the *buffer content distribution uniquely determines the variance function*, and it does so through the explicit formula:

$$V(t) \approx \inf_{\textsc{b} \geq 0} \frac{(\textsc{b} + (\textsc{c} - M)t)^2}{-2\log \mathbb{P}(Q > \textsc{b})}. \tag{7}$$

Hence, if we can estimate $\mathbb{P}(Q > \textsc{b})$, then the 'inversion formula' (7) can be used to retrieve the variance; notice that the infimum can be computed for any t, and consequently we get an approximation for the entire variance curve $V(\cdot)$ (of course up to some finite horizon).

Remark: Observe that in fact $\mathbb{P}(Q > \textsc{b})$ is also a function of the service speed $\textsc{c}$ (as the random variable Q depends on $\textsc{c}$). Interestingly, as the variance of the offered traffic in a window of length t, $V(t)$, does not depend on $\textsc{c}$, (7) entails that, if the approximation is correct, the minimum in the right-hand side should also not depend on $\textsc{c}$.

The unique correspondence between input and buffer content does not only apply for Gaussian processes, but can be found in many other situations as well. A well-known example is the $\textsc{m}/\textsc{g}/1$ queue. Assume Poisson arrivals stream (rate λ) of jobs, with service requirement distributed as random variable B (where its Laplace transform is denoted by $b(s) := \mathbb{E}e^{-sB}$), and service speed $\textsc{c}$. Let the load ϱ be defined as $\lambda \mathbb{E}B$. Denote by $q(\cdot; \textsc{c})$ the Laplace transform of the buffer content: $q(s; \textsc{c}) := \mathbb{E}e^{-sQ}$; the subscript '$\textsc{c}$' is added to emphasize the dependence on $\textsc{c}$. The Pollaczek-Khintchine formula says

$$q(s; \textsc{c}) = \left(1 - \frac{\varrho}{\textsc{c}}\right) \cdot \frac{s}{s - (\lambda/\textsc{c}) \cdot (1 - b(s))}.$$

This can be inverted to

$$b(s) = \frac{Cs}{\lambda q(s; C)} \cdot \left(\left(1 - \frac{\varrho}{C} \right) + \left(\frac{\lambda}{Cs} - 1 \right) q(s; C) \right).$$

We see similar phenomena as above: (i) the buffer-content distribution uniquely defines the distribution of the input process (the distribution of B); (ii) the left-hand side of the previous display clearly does not depend on C, so apparently the C in the right-hand side also cancels. $\diamond$

Example: Brownian bridge. In a Brownian bridge, time is restricted to the interval $[0, 1]$, and $V(t) = t(1 - t)$. It is easily verified that

$$\arg \inf_{t \in (0,1)} \frac{((C - M)t + B)^2}{2t(1 - t)} = \frac{B}{(C - M) + 2B};$$
$$I(B) = 2B(B + (C - M)).$$

We can now derive the upper bound on $V(t)$ from $I(B)$. To this end, we first compute

$$\arg \inf_{B \geq 0} \frac{((C - M)t + B)^2}{2I(B)} = \frac{(C - M)t}{1 - 2t};$$

then a lengthy calculation indeed gives $V(t) \leq t(1 - t)$, as desired. $\diamond$

B. Algorithm for estimating variance through inversion

In this section, we show how the inversion formula (7) can be used to estimate $V(\cdot)$. This inversion procedure consists of two steps. First we estimate the (complementary) buffer content distribution $\mathbb{P}(Q > \cdot)$, which is in the sequel abbreviated to BCD. Then we 'invert' the BCD to the variance curve $V(\cdot)$ by applying (7). We propose the following algorithm:

ALGORITHM: INVERSION.

1. Collect 'snapshots' of the buffer contents: $q_1, \ldots, q_N$; here q_i denotes the buffer content as measured at time $\tau_0 + i\tau$, for some $\tau > 0$. Estimate the BCD by the empirical distribution function of the q_i, i.e., estimate $\mathbb{P}(Q > B)$ by

$$\phi(B) := \frac{\#\{i : q_i > B\}}{N}.$$

2. Estimate $V(t)$ by

$$\inf_{B \geq 0} \frac{(B + (C - M)t)^2}{-2 \log \phi(B)},$$

for any $t \geq 0$. $\diamond$

Clearly, to obtain an accurate estimate of the BCD, both τ and N should be chosen sufficiently large. We come back to this issue in Section V.

C. Demonstration of inversion procedure

In the remainder of this section, we demonstrate the inversion approach through a simulation with synthetic (fBm) input; we emphasize, however, that the procedure could be performed for any other traffic process. We compare the estimated variance curve with the (known!) actual variance

Fig. 4. Sample BCD

curve. In this way we get a first impression of the accuracy of our approach; a more detailed numerical evaluation follows in Section V.

Consider a queue, in discrete time, with link rate C, fed by fBm. Its buffer dynamics are simulated as follows. A simulator [10] is used to generated fBm with a specific Hurst parameter $H \in (0, 1)$, yielding a list of numbers that represent the 'offered traffic per time slot'. These numbers serve as input to the queue: for every slot, the amount of offered traffic is added to the buffer content, while an amount equal to C is subtracted from the buffer content (where, when this number becomes negative, it is put to zero). Then, every $\tau \in \mathbb{N}$ slots, the queue's content is observed, yielding N snapshots $q_1, \ldots, q_N$ that are used to estimate $\mathbb{P}(Q > \cdot)$ (as in the above algorithm).

In this demonstration of the inversion procedure, we generate a fBm traffic trace with Hurst parameter $H = 0.7$; we take standard fBm, i.e., $M = 0$ and $V = t^{2H}$. The link capacity C is set to 0.8. The trace consists of 2^{24} slots, and we take snapshots of the buffer content every $\tau = 128$ slots.

We now discuss the output of the inversion procedure for our simulated example with fBm traffic. First, we estimate the BCD; a plot is given in Fig. 4. For presentation purposes, we plot the logarithm of the BCD, i.e., $\log \mathbb{P}(Q > \text{B})$. The BCD in Fig. 4 is 'less smooth' for larger values of B, which is due to the fact that large buffer levels are rarely exceeded, leading to less accurate estimates.

Second, we estimate the variance $V(t)$ for t equal to the powers of 2 ranging from 2^0 to 2^7, using the BCD, i.e., by using the algorithm. The resulting variance curve is shown in Fig. 5 ('inversion approach'). The minimization (over B) was done by straightforward numerical techniques. To get an impression of the accuracy of the inversion approach, we have also plotted the variance curve as can be estimated directly from the synthetic traffic trace (i.e., by using the 'direct approach' introduced earlier), as well as the real variance function for fBm traffic, i.e., $V(t) = t^{2H}$. The figure shows that the three variance curves are remarkably close to each other. This confirms that the inversion approach is an accurate way to estimate the burstiness. We note that the graph shows that the inversion approach slightly *overestimates* the variance. A more detailed validation of the

Fig. 5. Sample variance curves

inversion approach follows in the next section.

V. ERROR ANALYSIS OF THE INVERSION PROCEDURE

In the previous section the inversion approach was demonstrated, and it was shown to perform well for fBm with $H = 0.7$, under a specific choice of N and τ. Evidently, the key question is under what circumstances the procedure works. To this end, we first identify the three possible sources of errors:

- The inversion approach is based on the *approximation* (7).
- The BCD $\mathbb{P}(Q > \cdot)$ is *estimated*; there could still be an estimation error involved. In particular, one may wonder what the impact of the choice of N and τ is.
- The procedure *assumes* 'perfectly Gaussian' traffic, although real network traffic may not be (accurately described by) Gaussian, see Section III.

We will now quantitatively investigate the impact of each of these errors on our 'inversion approach'.

A. Approximation of the buffer content distribution.

In (6) an approximation of the BCD is given. As the inversion formula (7) is based on this approximation, evidently, errors in (6) might induce errors in the estimate of $V(\cdot)$. Therefore, we now assess the accuracy of (6). We focus on the practically relevant case of fBm; in line with the previous section, we choose $M = 0$ and $V(t) = t^{2H}$. Straightforward calculations now reveal that (6) implies

$$\log \mathbb{P}(Q > \text{B}) \approx -\frac{1}{2} \cdot \left(\frac{\text{B}}{(1-H)} \right)^{2-2H} \cdot \left(\frac{\text{C}}{H} \right)^{2H}.$$

We verify how accurate the approximation is, for two values of H: the Brownian case $H = 0.5$, and a case with long-range dependence $H = 0.7$ (which is a rather typical value, as found in many

Fig. 6. $\mathbb{P}(Q > \textsc{b})$ and theoretical approximation ($H = 0.5$)


Fig. 7. $\mathbb{P}(Q > \textsc{b})$ and theoretical approximation ($H = 0.7$)

measurement studies). Several runs of fBm traffic are generated (with different random seeds), with 2^{24} slots of traffic per run, which are used to simulate the buffer dynamics. For $H = 0.5$ we choose link rate $\textsc{c} = 0.2$, for $H = 0.7$ we choose $\textsc{c} = 0.8$; these choices of $\textsc{c}$ are such that the queue is non-empty sufficiently often (to make sure that a reliable estimate of the BCD is obtained). Figs. 6 and 7 show for the various runs the approximation of the BCD, as well as their theoretical counterpart. It can be seen that, in particular for small B, the empirically determined BCD matches very well with the values predicted by (6).

B. Estimation of the buffer content distribution.

The inversion formula requires the BCD $\mathbb{P}(Q > \cdot)$, which we approximate by $\phi(\cdot)$, i.e., its empirical counterpart. This may lead to errors; the impact of this error on the estimation of the

Fig. 8.   Comparing $\mathbb{P}(Q > \text{B})$ for various trace lengths

Fig. 9.   Comparing $\mathbb{P}(Q > \text{B})$ for various sampling intervals

variance curve is the subject of this subsection. It could be expected that the larger N (more observations) and τ (less correlation between the observations), the better the estimate.

We first investigate the impact of N. The simulator is run as in previous cases (with $H = 0.7$), with the difference that we only use the first $x\%$ of the snapshots samples to estimate the BCD. Figure 8 shows the estimation of the buffer content distribution, for various x ranging from 0.1% to 100%. The figure shows that, in particular for relatively small B, a relatively small number of observations suffices to obtain an accurate estimate.

Notice that we chose in our inversion procedure a *fixed* sampling frequency (i.e., τ^{-1}). It can be seen that this 'periodic sampling' is by no means necessary; the BCD-estimation procedure obviously still works when the sampling epochs are not equally spaced. In fact, one should realize that, if the sampling is performed in a purely periodic fashion, and if in addition (a substantial part of)

the traffic is also periodic, then one may obtain even unreliable estimates. Therefore, it may have advantages to sample at for instance Poisson epochs.

Second, we investigate the impact of the interval length between two consecutive snapshots τ. Figure 9 shows the determined BCD for τ ranging from observing every 32 to every 8192 slots. It can be seen that, particularly for small B the fit is quite good, even when the buffer content is polled only relatively rarely.

C. The impact of the Gaussianity assumption.

Approximation (6) explicitly assumes that the traffic process involved is Gaussian. We have seen in Section III that this claim is not always justified. Therefore, we now investigate how sensitive our inversion approach is with respect to the Gaussianity of the input traffic.

We study the impact of non-Gaussianity by mixing, for every slot, a fraction α of the generated fBm traffic with a fraction $1-\alpha$ traffic from an alternative (non-Gaussian) stream, before the mixture is fed into the queue. Note that the variance of the mixture is, in self-evident notation,

$$V(t) = \alpha^2 V_{[\text{fBm}]}(t) + (1 - \alpha)^2 V_{[\text{alt}]}(t).$$

We vary α from 1 to 0, to assess the impact of the non-Gaussianity.

The alternative input model that we choose here is the M/G/∞ input model (see Section III), inspired by, e.g., [1]–[3]. We denote the flow arrival rate by (λ). While in the system, traffic is generated at a constant rate r. In line with measurements studies (a classical reference is [8]), we choose Pareto(β) jobs, obeying the distribution function

$$F(x) = 1 - 1/(x + 1)^{\beta}, \quad \beta > 0.$$

As the objective is to assess the impact of varying the parameter α, we have chosen to select the parameters of the M/G/∞ input model such that it is 'compatible' with fBm, in that their means are equal, and their variances $V(\cdot)$ are 'similar' (in a sense that is defined below). This is achieved as follows.

- The means of both traffic stream are made compatible by adding a drift to the fBm inputs equal to the mean of the M/G/∞ traffic stream, i.e., $\lambda r/(\beta - 1)$. The Gaussianity of the fBm input is not affected by the addition of such a drift.
- To make the variances of both traffic streams 'compatible', we make use of a derivation in earlier work of the exact variance function $V(t)_{[\text{alt}]}$, see [20]; there it was shown that for the relevant case [8] of $\beta \in (1, 2)$, for large t,

$$V_{[\text{alt}]}(t) \approx 2r^2 \frac{\lambda t^{3-\beta}}{(3 - \beta)(2 - \beta)(\beta - 1)}.$$

  It is noted that long-range dependence is a property of long time-scales. We therefore chose to select the parameters of the M/G/∞ input model such that the ratio of the variances converges to a 1 as t goes ∞, thus guaranteeing that they possess the same 'degree of long-range dependence'.

Clearly, we can now estimate the remaining parameters and compute the variance of the traffic mixture.

Fig. 10. Variance curves for Gaussian/non-Gaussian traffic mixtures, $\alpha = 0.9$ and $\alpha = 0.8$

Fig. 11. Variance curves for Gaussian/non-Gaussian traffic mixture, $\alpha = 0$

The next step is to run, for different values of α, the simulation. We then determine the (theoretical) variance curve of the traffic mixture, and compare it to the variance curve found through the inversion approach. In Fig. 10 we focus on the 'nearly-Gaussian' cases $\alpha = 0.8$ and $\alpha = 0.9$, which are plotted together with their theoretical counterparts. The figure shows that the presence of non-Gaussian traffic has some, but no crucial impact on our inversion procedure.

We also consider the (extreme) case of $\alpha = 0$, i.e., no Gaussian traffic at all, to see if our inversion procedure still works. In Fig. 11 the various variance curves are shown: the theoretical curve, the curve based on the 'direct approach', as well as the curve based on the inversion approach. Although not a perfect fit, the curves look similar and still relatively close to each other (but, of course, the fit is worse than for $\alpha = 0.8$ and 0.9). Note that the non-Gaussian traffic may 'have some Gaussian characteristics', as argued in Section III.

We conclude that our simulation experiments show the 'robustness' of the inversion procedure. Despite the approximations involved, with a relatively low measurement effort, the variance curve is estimated accurately — even for traffic that is not 'perfectly Gaussian'. Given the evident advantages of the inversion approach over the 'direct approach' (minimal measurement effort required, retrieval of the entire variance curve $V(\cdot)$, etc. — see the discussion in Section III), the former method is to be preferred. In the next section we verify whether this conclusion also holds for real (i.e., not artificially generated) network traffic. This is done by inserting the estimated mean and variance into the dimensioning formulae, so that we can validate whether the resulting bandwidth values are such that the performance requirement is met.

VI. BANDWIDTH DIMENSIONING PROCEDURE, AND VALIDATION

As we have seen in Section V, the inversion approach shows rather good performance: even if the underlying traffic deviates from Gaussian, one still obtains a relatively good estimate of the variance function. It remains unclear however whether plugging in this variance function into the dimensioning formula, i.e., (4) or (5), also leads to good estimates of the required bandwidth. Another question is whether such conclusions remain valid for the traces from our data set (rather than artificially generated data). The purpose of this section is to study these issues. We begin by detailing our dimensioning approach.

A. Dimensioning procedure

We now describe our procedure to estimate the bandwidth needed in order to meet a predefined performance criterion. Suppose we have a trace of data, consisting of timestamps of packets, as well as the corresponding packet sizes; recall that in our data set each trace corresponds to 15 minutes (real time). We wonder what the minimum service rate is such that the performance requirement is satisfied.

ALGORITHM: DIMENSIONING.
1. Estimate the variance function $V(\cdot)$ by estimating the BCD and performing algorithm INVERSION.
2. Insert this estimated variance into the dimensioning formula, i.e., (4) or (5).   ◇

Notice that in the above procedure there was one choice left open: in the first step the BCD is determined by feeding the trace into the queue, and we did not specify the service speed, say c_q, of this queue. Clearly, c_q should not be chosen to high, because then the queue would be nearly always empty, leading to poor estimates of the BCD; we return to this issue later.

B. Validation

The idea behind this validation section is to systematically assess the accuracy of our dimensioning approach. We do so by decoupling two effects.
- *Validation of the required bandwidth formula*. Here the question is: suppose we are given perfect information about the variance function, and we plug this into our dimensioning formula, how good is the resulting estimate for the required bandwidth?

- *Impact of estimation errors in $V(\cdot)$ on required bandwidth.* Here test how the errors in the estimate of $V(\cdot)$, caused by the inversion procedure, have impact on the estimate of the required bandwidth.

The 'decoupling' allows us to gain precise insight into — and thus a proper validation of — both steps of the above link dimensioning procedure. We focus on the criterion of link transparency, and the corresponding dimensioning formula (4); with the same methodology, one can perform the validation with respect to the buffer overflow criterion.

Validation of the required bandwidth formula: We now check the accuracy of the dimensioning formula (4). It requires knowledge of M and $V(T)$. As we do not have their 'real values', we estimate them using the 'direct approach'. We emphasize that this direct approach has significant disadvantages in practice, see the discussion in Section IV. However, in order to assess the accuracy of (4), we do not have any alternative.

In more detail: from each trace we estimate the average traffic rate M and the variance $V(T)$ of the offered traffic at timescale T; with A_i denoting the amount of traffic offered over the i-th interval of length T,

$$\widehat{M} = \frac{1}{nT} \sum_{i=1}^{n} A_i, \quad \text{and} \quad \widehat{V}(T) = \frac{1}{n-1} \sum_{i=1}^{n} (A_i - M)^2.$$

Then the resulting estimates, as well as the specified values of T and ε, are inserted into (4) to obtain the (estimated) required bandwidth.

We choose to determine the average traffic rate M per 15 minutes (recall that each trace contains 15 minutes of traffic), and set T to 1 sec., 500 msec. and 100 msec. (and thus determine the variance at those timescales), which are, for various applications, timescales that are important to the perception of quality by (human) users. We set ε to 1%. Importantly, note that these settings for T and ε are just examples — network providers can choose the setting that suits their (business) needs best.

In order to validate if the estimated bandwidth capacity C indeed corresponds to the required bandwidth, we introduce the notion of 'realized exceedance', denoted with $\hat{\varepsilon}$. We define the 'realized exceedance' as the fraction of intervals of length T, in which the amount of offered traffic A_i exceeds the estimated required capacity CT — we stress the fact that 'exceedance' in this context does not correspond to 'packet loss'. In other words:

$$\hat{\varepsilon} := \frac{\#\{A_i \mid A_i > \text{C}T\}}{n} \quad (i \in 1 \ldots n).$$

If C is properly dimensioned, then 'exceedance' (as in $A(T) > \text{C}T$) may be expected in a fraction ε of all intervals. There are, however, (at least) two reasons why $\hat{\varepsilon}$ and ε may not be equal in practice. (i) Firstly, (4) assumes 'perfectly Gaussian' traffic, which is, as we have seen, not always the case. Evidently, deviations of 'perfectly Gaussian' traffic may have an impact on the estimated C. (ii) Secondly, to obtain (1), an upper bound (viz. the Chernoff bound) on the target probability has been used, and it is not clear upfront how far off this bound is. To assess to what extent the dimensioning formula for Gaussian traffic is accurate for real traffic, we compare ε and $\hat{\varepsilon}$. We do this comparison for the hundreds of traces that we collected at measurement locations $\{\text{U}, \text{R}, \text{C}, \text{A}, \text{S}\}$.

Table I presents the average differences between the targeted ε and the 'realized exceedance' $\hat{\varepsilon}$ at each location, as well as the standard deviations, for three different timescales T. The table shows

Location	T	avg $\lvert \varepsilon - \hat{\varepsilon} \rvert$	stderr $\lvert \varepsilon - \hat{\varepsilon} \rvert$	$\bar{d}$	σ_d
U	1 sec	0.0095	0.0067	1.33	0.10
	500 msec	0.0089	0.0067	1.35	0.09
	100 msec	0.0077	0.0047	1.42	0.09
R	1 sec	0.0062	0.0060	2.91	1.51
	500 msec	0.0063	0.0064	3.12	1.57
	100 msec	0.0050	0.0053	3.82	1.84
C	1 sec	0.0069	0.0047	1.71	0.44
	500 msec	0.0066	0.0043	1.83	0.49
	100 msec	0.0055	0.0041	2.13	0.67
A	1 sec	0.0083	0.0027	1.13	0.03
	500 msec	0.0083	0.0024	1.14	0.03
	100 msec	0.0079	0.0020	1.19	0.03
S	1 sec	0.0052	0.0050	1.98	0.78
	500 msec	0.0049	0.0055	2.10	0.87
	100 msec	0.0040	0.0059	2.44	1.01

TABLE I

REQUIRED BANDWIDTH: ESTIMATION ERRORS AND DIMENSIONING FACTOR ($\varepsilon = 0.01$).

that differences between ε and $\hat{\varepsilon}$ are small. Hence, the required bandwidth to meet a pre-specified performance target of ε is estimated accurately.

When dimensioning a network link, network providers often use 'rules of thumb', such as C $= \alpha\, M$, for a given number α (for instance: take the mean rate, increased by 50%). To verify whether such a (simplistic) approach could work, it is interesting to get an idea of the required 'dimensioning factor', i.e., the (estimated) required bandwidth capacity compared to the average load (i.e., C$/M$). These dimensioning factors and their standard deviations, averaged over all traces at each location, are given in the two rightmost columns of Table I, for $T = 1$ sec., 500 msec., and 100 msec. It shows, for instance, that at location U, some 33% extra bandwidth capacity would be needed on top of the average traffic load M, to cater for 99% ($\varepsilon = 0.01$) of all traffic peaks at a timescale of $T = 1$ sec. At location R, relatively more extra bandwidth is required to meet the same performance criterion: about 191%. Such differences between those locations can be explained by looking at the network environment: at location R, a single user can significantly influence the aggregated traffic, because of the relative low aggregation level (tens of concurrent users) and the high access link speeds (100 Mbit/s, with a 1 Gbit/s backbone); at location U, the user aggregation level is much higher, and hence, the traffic aggregate is 'more smooth'. Conclusion is that simplistic dimensioning rules of the type C $= \alpha\, M$ are inaccurate, as the α is all but a universal constant (it depends on the nature of the traffic, on the level of aggregation, the network infrastructure, and on the performance target imposed).

The dimensioning factors (cf. Table I) for the present case studies can be obtained as the ratio of C and M, at certain T and ε. As indicated above, the dimensioning factor increases when the performance criterion (through ε and T) becomes more stringent. To give a few examples of the impact of the performance parameters T and ε on the required bandwidth capacity, we plot curves
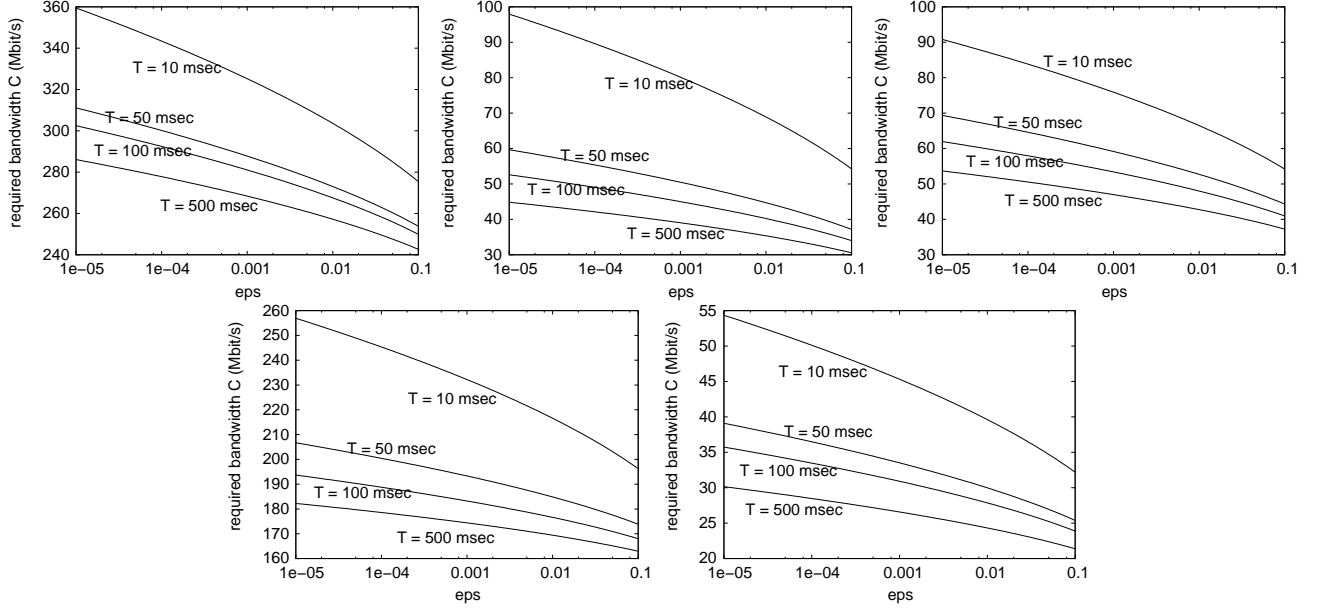
Fig. 12. Required bandwidth for other settings of T and ε for locations $\{\mathrm{U,R,C,A,S}\}$

for the requires bandwidth capacity at $T = 10, 50, 100$ and 500 msec., and ε ranging from 0.00001 to 0.1, in Fig. 12. In these curves, M and $V(T)$ are (directly) estimated from an example traffic trace collected at each of the locations $\{\mathrm{U,R,C,A,S}\}$.

Figure 12 shows that the required bandwidth C decreases in both T and ε, which is intuitively clear. The figures show that C is more sensitive to T than to ε — take for instance the top-left plot in Figure 12, i.e., location U, example trace #1; at $\varepsilon = 10^{-5}$, the difference in required bandwidth between $T = 10$ msec. and $T = 100$ msec., is some 20%. At $T = 100$ msec., the difference in required bandwidth between $\varepsilon = 10^{-5}$ and $\varepsilon = 10^{-4}$ is just 3% approximately.

We have verified whether the required bandwidth is accurately estimated for these case-studies with different settings of T and ε. The estimation errors in these new situations are similar to the earlier obtained results (cf. Table I). It should be noted however, that we have not been able to verify this for all possible combinations of T and ε: for $\varepsilon = 10^{-5}$ and $T = 500$ msec for instance, there are only 1800 samples in our traffic trace (which has a length of 15 minutes) and hence, we cannot compute the accuracy of our estimation. Another remark that should be made here, is that for locations with only limited aggregation in terms of users (say some tens concurrent users), combined with a small timescale of $T = 10$ msec, the traffic is no longer Gaussian (i.e., $\gamma \ll 1$). Consequently, the accuracy of our required bandwidth estimation decreases.

Impact of estimation errors in the variance on the required bandwidth: In Section V we have seen for artificially generated traffic that the inversion worked well, and the first part of the present section assessed the quality of the bandwidth dimensioning formula (4). We now combine these two elements: we perform the inversion, and we insert the resulting variance into (4), and see how well this predicts the required bandwidth.

First, we estimate the average traffic rate M and variance $V(T)$ directly from the offered traffic stream, as above. Also, we 'replayed' our traces by inserting them into a virtual queue with capacity C_q (emulated in a `Perl` script), so that we can use the inversion algorithm to estimate $V(T)$. We

denote this directly estimated variance by $V_{\text{direct}}(T)$, and the estimate resulting from inversion by $V_{\text{inversion}}(T)$. Again, we emphasize that the direct estimator $V_{\text{direct}}(T)$ may be infeasible in operational environments, as T is likely chosen to be rather small. In our case studies, we have chosen to set the buffer occupancy sampling interval to 1 sec., which ensures that we have a sufficient number of snapshots to reliably estimate the BCD (see Section V). Furthermore, we set T, the timescale that we aim to determine the variance for, to 100 msec.

The remaining parameter to be set is C_q, the queue's service rate. Clearly, when C_q is chosen too small, say $C_q < M$, the system is not stable in that it cannot serve all the traffic offered. Hence, we should have $C_q \geq M$. On the other hand, if C_q is much larger than M, then the queue's occupancy will, obviously, be zero at most observation times. This would lead to an unreliable estimation of the BCD, so C_q should not be too large.

We have performed numerous experiments to see if there is a general guideline for choosing C_q, for instance $C_q = c \cdot M$ where c is some constant, that ultimately leads to an accurate approximation of $V_{\text{inversion}}(T)$. As it turns out, there is no single value of c or C_q that works well in all cases. Choosing C_q anywhere between $1 \cdot M$ and $2 \cdot M$, however, gives reasonable results: $V_{\text{inversion}}(T)$ is always in the same order of magnitude as $V_{\text{direct}}(T)$, and in most cases within 5-20% of each other. Furthermore, we stress that the impact of an estimation error in $V_{\text{inversion}}(T)$ on our ultimate goal of bandwidth dimensioning is somewhat mitigated: in (4) first the square-root of the variance is taken, and then it is added to some average rate M. To study the accuracy of the estimation we introduce

$$\nu := \frac{V_{\text{inversion}}(T)}{V_{\text{direct}}(T)}.$$

As we are interested in using the variance estimates for dimensioning purposes, we compare the required bandwidth C_{direct} (i.e., bandwidth estimated using the directly estimated variance) and $C_{\text{inversion}}$ (i.e., bandwidth estimated using the inversion approach). Recall that C_{direct} is an accurate estimation of the required bandwidth, as we saw earlier in the present subsection. We also compare $C_{\text{inversion}}$ and the 'empirical' minimally required bandwidth $C_{\text{empirical}}$:

$$C_{\text{empirical}} := \min\left\{ C \ : \ \frac{\#\{A_i \mid A_i > CT\}}{n} \leq \varepsilon \right\},$$

to be interpreted as the minimum bandwidth that provides the trace with the desired performance. Also, we introduce Δ as an indicator of the quality of the estimation of the required bandwidth through the inversion approach, with respect to both C_{direct} and $C_{\text{empirical}}$:

$$\Delta_{\text{var}} := \frac{C_{\text{inversion}}}{C_{\text{direct}}}, \quad \text{and} \quad \Delta_{\text{cap}} := \frac{C_{\text{inversion}}}{C_{\text{empirical}}}.$$

Table II lists the validation results of a number of 'test traces': it shows that the variances are rather accurately estimated. Table III compares the estimated required capacity (with $\varepsilon = 0.01$, $T = 100$ msec.) computed via both the direct and inversion approaches. Also the empirically found minimum required bandwidth is tabulated.

As can immediately be seen from the values of Δ_{var}, the required bandwidth capacity as estimated through the inversion approach to estimate the variance, is remarkably close to that obtained through the direct approach: on average, the differences are less than 1%. Also, comparison with

trace	M	$V_{\mathrm{direct}}(T)$	$V_{\mathrm{inversion}}(T)$	ν
loc. U #1	207.494	3.773	4.026	1.067
#2	238.773	7.313	7.690	1.052
loc. R #1	18.927	0.492	0.483	0.981
#2	3.253	0.058	0.062	1.062
loc. C #1	23.894	0.633	0.644	1.018
#2	162.404	10.650	12.379	1.162
loc. A #1	147.180	0.939	1.064	1.133
#2	147.984	0.745	0.747	1.003
loc. S #1	14.254	0.200	0.201	1.004
#2	2.890	0.023	0.023	1.022

TABLE II

VALIDATION RESULTS FOR THE BURSTINESS ESTIMATION METHODOLOGY (M IS IN MBIT/SEC, V IS IN MBIT^2).

trace	$\mathrm{C}_{\mathrm{empirical}}$	$\mathrm{C}_{\mathrm{direct}}$	$\mathrm{C}_{\mathrm{inversion}}$	Δ_{var}	Δ_{cap}
loc. U #1	258.398	266.440	268.385	1.007	1.039
#2	302.663	320.842	322.934	1.007	1.067
loc. R #1	37.653	40.221	40.020	0.995	1.063
#2	10.452	10.568	10.793	1.021	1.033
loc. C #1	44.784	48.033	48.250	1.005	1.077
#2	265.087	261.444	269.182	1.030	1.015
loc. A #1	171.191	176.588	178.480	1.011	1.043
#2	168.005	174.178	174.218	1.000	1.037
loc. S #1	27.894	27.843	27.873	1.001	0.999
#2	7.674	7.482	7.532	1.007	0.981

TABLE III

VALIDATION RESULTS FOR THE BURSTINESS ESTIMATION METHODOLOGY (CONTINUED) (C IS IN MBIT/SEC).

the empirical minimum required bandwidth, through Δ_{cap}, shows that the use of the inversion procedure leads to estimates for the required bandwidth that are, remarkably, on average less than 4% off. Comparing the respective values for ν and Δ_{var} in Table II and Table III, one observes that an estimation error in $V(T)$, indeed, has only limited impact on the error in C.

Now we have seen the impact of using indirectly estimated variances for a number of test traces (i.e., through the inversion approach), it remains to assess the overall accuracy of our dimensioning approach, we have computed the Δ values as described above for *all* our traces at every location — see the upper part of Table IV. We have tabulated the average values of Δ_{var} and Δ_{cap}, as well as their standard error terms. The required bandwidth estimations are remarkably accurate. In the lower part of Table IV, we have tabulated the same metrics, but only used the traces that are 'fairly Gaussian', in that their linear correlation coefficient is above 0.9. This improves the results even further — leading to the conclusion that errors in the required bandwidth estimation using the inversion approach to estimate the variance, are primarily caused by non-Gaussianity of the offered traffic.

location	avg Δ_{var}	stderr Δ_{var}	avg Δ_{cap}	stderr Δ_{cap}
U	1.00	0.01	1.01	0.07
R	0.96	0.10	0.90	0.19
C	1.00	0.03	1.04	0.11
A	1.00	0.01	1.04	0.02
S	1.00	0.03	0.99	0.10

location	avg Δ_{var}	stderr Δ_{var}	avg Δ_{cap}	stderr Δ_{cap}
U	1.00	0.01	1.03	0.06
R	1.00	0.02	1.00	0.10
C	1.00	0.02	1.05	0.08
A	1.00	0.01	1.04	0.01
S	1.00	0.01	1.01	0.05

TABLE IV

VALIDATION RESULTS FOR THE BURSTINESS ESTIMATION METHODOLOGY (OVERALL RESULTS) — UPPER TABLE: ALL TRACES; LOWER TABLE: TRACES WITH $\gamma > 0.9$.

VII. DISCUSSION AND CONCLUDING REMARKS

This paper provides a framework for link dimensioning. We first derived generic dimensioning formulae, which require knowledge of the variance of the input traffic. Then we proposed an efficient variance estimation technique. This so-called inversion method was extensively tested in experiments with real network traffic from various representative networking environments. These tests showed that the resulting estimated link rates are remarkably accurate.

In the remainder of this section we discuss a number of directions for further research, and we reflect on alternative approaches.

Extensions, further research. In this paper we focused on dimensioning of a single network link. Clearly, dimensioning may get more complicated when (optimizing the) dimensioning (of) an entire network: How to guarantee a performance target network-wide? Also, other factors may come into play, e.g., routing can also help in improving network-wide performance.

A prerequisite for our method is the *performance criterion*; when adopting the link transparency criterion, one needs to specify appropriate values for the parameters T and ε. These values are dictated by the so-called *perceived Quality-of-Service*, i.e., the performance as perceived by users. Clearly any application has other requirements, but these are also not necessarily uniform among users of the same application — performance is a subjective issue. For traditional services (such as voice) the mapping from (subjective) perceived performance to (objective) performance parameters (such as T and ε) has been intensively studied; for more advanced applications (for instance interactive web applications) this is still an open issue.

Often, dimensioning issues cannot be addressed without taking into account pricing aspects. Put differently, the utility a network user assigns to the service he is offered depends not only on the performance level, but also on the price he has to pay for it; in fact, the combination of performance and price will determine whether a potential user is willing to subscribe. From this perspective it is not the providers' task to choose the minimum bandwidth such that a pre-specified performance target is met, but rather to choose bandwidth and prices such that some objective function (for

instance revenues, or profits) is optimized.

Other approaches. The purpose of our inversion method is to retrieve the essential traffic characteristics at low measurement costs. We remark that several other 'cheap' (i.e., with low measurement effort) methods have been proposed in the literature. We now discuss some of these.

If one cannot assume Gaussianity, one could still use approximation (2), and estimate the mgf $\mathbb{E}e^{\vartheta A(-t,0)}$ (as a function of both t and ϑ) from traffic measurements. There are several papers on its statistical aspects, see for instance [14]. As we have seen, the formula simplifies considerably when assuming Gaussian traffic, as then estimating the mgf reduces to estimating M and $V(\cdot)$.

The paper by Duffield *et al.* [11] presents a procedure to accurately estimate the *asymptotic cumulant function*

$$\Lambda(\vartheta) := \lim_{t \to \infty} \frac{1}{t} \log \mathbb{E}e^{\vartheta A(t)}$$

from traffic measurements — this function is useful, because, under some assumptions on the traffic arrival process, it holds that $\mathbb{P}(Q > \text{B}) \approx \exp(-\vartheta^\star \text{B})$, for B large, where $\vartheta^\star$ solves the equation $\Lambda(\vartheta) = \text{C}\vartheta$. The crucial assumption, however, is that the arrival process be short-range dependent; otherwise one cannot be sure that the cumulant exist —think of fBm, for which $\mathbb{E}\exp(\vartheta A(t))$ is of the form $\exp(M\vartheta t + \frac{1}{2}\vartheta^2 t^{2H})$, so that, for $H \in (\frac{1}{2}, 1)$,

$$\lim_{t \to \infty} \frac{1}{t} \log \mathbb{E}e^{\vartheta A(t)} = \lim_{t \to \infty} \left(\vartheta M + \frac{1}{2}\vartheta^2 t^{2H-1} \right) = \infty.$$

The fact that this method cannot be used for long-range dependent traffic makes its use for dimensioning purposes limited. We remark that a crucial difference with our approach is that [11], [14] measure *traffic*, whereas we propose to measure (or, better: to sample) the *buffer content*.

Another related study is by Kesidis *et al.* [7]. Like in our method, their approach relies on the estimation of the buffer content distribution $\log \mathbb{P}(Q > \text{B})$. Under the assumption of short-range dependent input, $\log \mathbb{P}(Q > \text{B})$ is linear for large B (with slope $-\vartheta^\star$). Having estimated $\vartheta^\star$ the probability of overflow over higher buffer levels can be estimated, by extrapolating $\log \mathbb{P}(Q > \text{B})$ linearly. For long-range dependent traffic $\log \mathbb{P}(Q > \text{B})$ is not linear (see Fig. 7), and hence this method cannot be used.

Another approach to minimize the measurement effort was presented in Van den Berg *et al.* [5]. There traffic is assumed to be Gaussian, but with the correlation structure of the M/G/∞ input model. Interestingly, under the link transparency criterion the required bandwidth simplifies to $\text{C} = M + \alpha\sqrt{M}$, where α depends on the performance criterion (i.e., T and ε), and on the characteristics of a single flow (i.e., the distribution of the flow duration B and the traffic rate r), but not on the traffic arrival rate λ. This property enables a simple estimate of the additionally required bandwidth if in a future scenario traffic growth is mainly due to a change in λ (e.g., due to growth of the number of subscribers), and not due to changes in user behavior.

VIII. Acknowledgment

REFERENCES

[1] R. Addie, P. Mannersalo, and I. Norros (2002). Most probable paths and performance formulae for buffers with Gaussian input traffic. *European Transactions on Telecommunications*, Vol. 13, pp. 183-196.

[2] C. Barakat, P. Thiran, G. Iannaccone, C. Diot, and Ph. Owezarski (2003). Modeling Internet backbone traffic at the flow level. *IEEE Transactions on Signal Processing*, Vol. 51, pp. 2111-2114.

[3] N. Ben Azzouna, F. Clérot, C. Fricker, and F. Guillemin (2004). A flow-based approach to modeling ADSL traffic on an IP backbone link. To appear in *Annals of Telecommunications*.

[4] J. Beran (1994). *Statistics for Long-Memory Processes*, Chapman & Hall/CRC.

[5] H. van den Berg, M. Mandjes, R. van de Meent, A. Pras, F. Roijers, and P. Venemans (2006). QoS-aware bandwidth provisioning of IP links. *Computer Networks*, Vol. 50, pp. 631-647.

[6] D. Botvich and N. Duffield (1995). Large deviations, the shape of the loss curve, and economies of scale in large multiplexers. *Queueing Systems*, Vol. 20, pp. 293-320.

[7] C. Courcoubetis, G. Kesidis, A. Ridder, J. Walrand, and R. Weber (1995). Admission control and routing in ATM networks using inferences from measured buffered occupancy. *IEEE Transactions on Communications*, Vol. 43, pp. 1778–1784.

[8] M. Crovella, A. Bestavros (1997). Self-similarity in World Wide Web traffic. Evidence and possible causes. *IEEE/ACM Transactions on Networking*, Vol. 6, pp. 835–846.

[9] K. Dębicki and Z. Palmowski (1999). Heavy-traffic Gaussian asymptotics of on-off fluid model. *Queueing Systems*, Vol. 33, pp. 327–338.

[10] T. Dieker (2002). Fractional Brownian motion simulator. Available at `http://homepages.cwi.nl/~ton/fbm/index.html`.

[11] N. Duffield, J. Lewis, N. O'Connell, R. Russell, and F. Toomey (1995). Entropy of ATM traffic streams: a tool for estimating QoS parameters. *IEEE Journal on Selected Areas in Communications*, Vol. 13, pp. 981-990.

[12] C. Fraleigh, F. Tobagi, and C. Diot (2003). Provisioning IP backbone networks to support latency sensitive traffic. *Proceedings IEEE Infocom 2003*, San Francisco CA, USA.

[13] A. Ganesh, N. O'Connell, and D. Wischik (2004). *Big Queues.* Springer Verlag, Berlin.

[14] L. Györfi, A. Rácz, K. Duffy, J. Lewis, and F. Toomey (2000). Distribution-free confidence intervals for measurements of effective bandwidth. *Journal of Applied Probability*, Vol. 37, pp. 224–235.

[15] I. Juva, R. Susitaival, M. Peuhkuri, and S. Aalto (2005). Traffic characterization for traffic engineering purposes: analysis of Funet data. *Proceedings 1st EURO-NGI Conference*, Rome, Italy.

[16] F. Kelly (1996). Notes on effective bandwidths. In: *Stochastic Networks: Theory and Applications* (Editors F.P. Kelly, S. Zachary and I.B. Ziedins). Oxford University Press, pp. 141-168.

[17] J. Kilpi and I. Norros (2002). Testing the Gaussian approximation of aggregate traffic. *Proceedings Internet Measurement Workshop*, Marseille, France. Available at `http://www.vtt.fi/tte/rd/traffic-theory/papers/`

[18] W. Leland, M. Taqqu, W. Willinger, and D. Wilson (1994). On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, Vol. 2, pp. 1-15.

[19] M. Mandjes and R. van de Meent (2005). Inferring traffic burstiness by sampling the buffer occupancy. In: R. Boutaba, K. Almeroth, R. Puigjaner, S. Shen, J. Black (eds.): *Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; Mobile and Wireless Communication Systems.* Fourth International IFIP-TC6 Networking Conference (Networking 2005), Waterloo, Canada. Lecture Notes in Computer Science (LNCS) Series, 3462, pp. 303-315.

[20] M. Mandjes, I. Saniee, and A. Stolyar (2005). Load characterization, overload prediction, and load anomaly detection for voice over IP traffic. *IEEE Transactions on Neural Networks*, Vol. 16, pp. 1019-1028.

[21] R. van de Meent and M. Mandjes (2005). Evaluation of 'user-oriented' and 'black box' traffic models for link provisioning. *Proceedings 1st EURO-NGI Conference*, Rome, Italy.

[22] R. van de Meent, M. Mandjes, and A. Pras (2006). Gaussian traffic everywhere? *IEEE International Conference on Communications*, Istanbul, Turkey.

[23] R. van de Meent (2006). Network link dimensioning: a measururement & modeling based approach. *Ph.D. thesis, University of Twente*.

[24] I. Norros (1994). A storage model with self-similar input. *Queueing Systems*, Vol. 16, pp. 387-396.

[25] I. Norros (1995). On the use of fractional Brownian motion in the theory of connectionless networks. *IEEE Journal on Selected Areas in Communications*, Vol. 13, pp. 953-962.

[26] V. Paxson and S. Floyd (1995). Wide-area traffic: the failure of Poisson modeling. *IEEE/ACM Transactions on Networking*, Vol. 3, pp. 226-244.

[27] G. Seres, A. Szlávik, J. Zátonyi, J. Bíró (2002). Alternative admission rules based on the many-sources asymptotics. *Proceedings Seventh International Symposium on Computers and Communications (ISCC'02)*, pp. 995-1000.

Michel Mandjes received M.Sc. (in both mathematics and econometrics) and Ph.D. degrees from the Vrije Universiteit (VU), Amsterdam, the Netherlands. After having worked as a member of technical staff at KPN Research (Leidschendam, the Netherlands) and Bell Laboratories/Lucent Technologies (Murray Hill NJ, USA), as a part-time full professor at the University of Twente, and as department head at the Centre for Mathematics and Computer Science (CWI), Amsterdam, he currently holds a full professorship at the University of Amsterdam, the Netherlands. His research interests include performance analysis of communication networks, queueing theory, Gaussian traffic models, traffic management and control, and pricing in multi-service networks.

Remco van de Meent received M.Sc (2001, in computer science) and Ph.D. (2006, for his thesis 'Network link dimensioning: a measurement and modeling based approach') degrees from the University of Twente, Enschede, the Netherlands. He is currently working as researcher at Virtu, Enschede, the Netherlands, and as post-doct at the University of Twente. His research interests include network dimensioning, traffic modeling, network security and network management.