Evaluation of ‘user-oriented’ and ‘black-box’ traffic models for link provisioning

Remco van de Meent
University of Twente
r.vandemeent@utwente.nl

Michel Mandjes
CWI
michel@cwi.nl

Abstract—To offer users a sufficient performance level, network links should be properly provisioned. The required bandwidth capacity may be determined through the use of a model of the real network traffic. In this paper, we study the use of two classes of traffic models: (i) ‘user-oriented models’, which capture the behavior of individual flows, and (ii) ‘black-box models’, which statistically describe the *superposition* of many users (and do not distinguish between individual flows). User-oriented models have the advantage that they allow for sensitivity analysis: the impact of a change in the user parameters (access rate, flow-size distribution) can be assessed. In general, however, our measurements indicated that black-box models are easier to estimate, and yield accurate provisioning guidelines.

I. INTRODUCTION

Traditionally, bandwidth provisioning in IP backbones is done by applying a set of ‘rules of thumb’. For instance, a commonly used procedure is to collect 5-minute measurements (e.g., by using SNMP) of the traffic offered, take its, say, 95% quantile, and add some safety margin. As these methods do not use the detailed characteristics of the traffic processes involved (i.e., fine-grained rate fluctuations), it is evident that such provisioning procedures cannot provide any performance guarantees that relate to these detailed timescales. Hence, as the user-perceived performance is crucially affected by the availability of sufficient resources over relatively detailed timescales, it is clear that the ‘coarse’ procedure described above has serious drawbacks.

Consequently, in order to do adequate provisioning (i.e., not too little bandwidth capacity, as that would negatively affect the performance; and also not too much, as that would evidently lead to a waste of resources), it is necessary to have a more thorough understanding of the detailed statistical properties of the traffic offered; information from coarse measurements does not suffice. Clearly, link provisioning would greatly benefit from the availability of accurate traffic models, as these would facilitate the prediction of the offered performance as a function of the amount of allocated resources (i.e., link bandwidth).

There are different classes of traffic models, with their specific pros and cons. Recently, much attention has been paid to ‘flow-level models’. These group the aggregate traffic streams into ‘flows’, which are ‘coherent strings’

of packets, for instance packets within the same TCP connection or UDP stream, or packets with the same origin-destination pair, etc. Flows arrive according to some random process (usually one assumes a Poisson process), stay in the system for some random time, and during their stay they transmit packets (for instance at a constant rate). We could call such a model ‘user-oriented’, as it relates to the traffic profiles of the users.

The most prominent advantage of using this type of models is that it facilitates sensitivity analysis. For instance, it enables the assessment of the effect of the migration of (a part of) the user population from a ‘slow’ access technology to a ‘faster’ one: what is the impact on the bandwidth needed? Also the effect of a change in the flow-size distribution could be quantified.

A second class of models could be named ‘black-box models’: they do not model individual flows, but rather attempt to find an accurate statistical description of the *aggregate* of all users. A commonly used subclass of black-box models are the *Gaussian* models. With $A(s, t)$ denoting the amount of traffic arriving in $[s, t)$, a Gaussian model with stationary increments is such that $A(s, t)$ only depends on the interval length $t - s$. More specifically, $A(s, t)$ follows a Normal law, with mean $\mu(t - s)$ (for some mean $\mu > 0$) and variance $v(t - s)$ (for some nonnegative function $v(\cdot)$), for any s, t such that $s \leq t$.

It is clear that the Gaussian model is in some sense an ‘artifact’, as it, at least in principle, allows for negative input. However, when $\mu(t - s)$ is substantially larger than the standard deviation $\sqrt{v(t - s)}$ it is highly unlikely that over an interval of length $t - s$ the increment is negative.

Both modeling approaches have their advantages and disadvantages, and this is exactly what we want to assess in this paper. Let us consider the following aspects:

- Black-box models abstract from the relation with the ‘user-level’, in that they only model the superposition of (usually many) flows. An immediate consequence is that this type of modeling does not lend itself to performing sensitivity analysis with respect to the ‘user-level parameters’ (such as access rates, etc.).
- Usually there is a strong heterogeneity between flows. This can have several causes. In the first place, the end users use different applications, which are character-

ized by different bandwidth consumption patterns. For instance: streaming applications could use a constant bit rate (possibly well below the access rate), whereas file transfers are based on TCP (and grab as much bandwidth as possible, constrained by the access rate, the maximum window size, and the bottleneck elsewhere in the network). A second (perhaps more important) cause of heterogeneity lies in the fact that the bottleneck for different flows could be at different links or routers somewhere else in the network. For instance, two downloads from different servers at different locations in the network, could result in very different transmission rates.

- Flow-level models use the notion of flow. Flows could be defined as TCP connections or UDP streams, but these could be small and numerous. A practical alternative is to do some aggregation, for instance by identifying flows as transmissions between origin-destination pairs, as long as the ‘gaps’ between packets do not exceed some predefined interval τ . Of course, the parameter τ is a ‘tuning knob’ for which an appropriate value needs to be selected.

The heterogeneity described above makes it worthwhile to split the user population in several subclasses, with their own characteristics. For instance, flows with a size (in bits) smaller than f could behave very differently from call bigger than f (cf. mice and elephants). But of course, then this parameter f should also be chosen and tuned, so this leads to a similar problem.

In addition, when the set of applications changes (which happens every now and then), these parameters have to be tuned again.

Related work. There is a vast body of literature on network traffic models; we mention a few studies that are particularly relevant in the scope of the present study. A classical contribution is by Leland *et al.* [1], who establish the self-similar nature of Ethernet-traffic; also for other types of traffic self-similar models, such as *fractional Brownian motion* (fBm), have proved to be adequate. These models are typically of the black-box type, as they model the aggregate of a (large) number of users. The use of fBm was further motivated by showing that it arises as a limiting model of a specific user-oriented model: the superposition of many on-off flows with heavy-tailed on- or off-times converges (after rescaling time) to fBm, see Crovella and Bestavros [2]. For smaller timescales, it is noted that Gaussianity cannot be assumed, see for instance Kilpi and Norros [3] and Fraleigh *et al.* [4]. Recently attention shifted somewhat to user-oriented models, such as the flow-level models of Ben Fredj *et al.* [5] and Barakat *et al.* [6]. We also mention a recent study by Ben Azzouna *et al.* [7] who succeed in finding a detailed flow-level description of ADSL traffic by decomposing the aggregate stream into

several classes, of which the parameters are estimated separately. As indicated above, the thresholds characterizing the different ‘flow groups’ are parameters, and need to be tuned; it is also noted that some components are estimated by black-box models (i.e., models in which the user behavior cannot be recognized).

Contribution. In this paper we develop bandwidth provisioning formulae for various traffic modeling approaches, both of user-oriented and black-box nature. These models are systematically compared, considering different scenarios (i.e., different access technologies, different aggregation levels), to investigate which type of modeling is more appropriate. Our analysis shows that, particularly due to the strong heterogeneity, flow-level models (of the type $M/G/\infty$) are often inappropriate. Gaussian models do fit nicely, but have the inherent drawback that they hardly allow for sensitivity analysis. We also show the accuracy of the resulting provisioning guidelines.

Approach and organization. In Section 2 we treat a number of preliminaries; in particular we introduce two important traffic models, namely $M/G/\infty$ traffic (which is a user-oriented model) and Gaussian traffic (which is a black-box model). In Section 3, we attempt to fit the $M/G/\infty$ model for various traces, and find a number of intrinsic difficulties. Section 4 shows that Gaussian models are easy to apply, and lead to accurate provisioning guidelines. Section 5 concludes.

II. SOME PRELIMINARIES ON LINK PROVISIONING AND TRAFFIC MODELING

Provisioning of network resources addresses the interrelationship between (i) offered traffic (in terms of both average load and burstiness), (ii) desired level of performance, and (iii) the required capacity. Generally, more capacity is needed when offered load and burstiness increase, or when the performance criterion becomes more stringent. To operate a network in a viable way, provisioning procedures balancing (i), (ii) and (iii) are required: scarce provisioning inevitably leads to performance degradation, whereas (too much) over-provisioning results in a waste of resources.

In the present paper, we model traffic with the following (performance) objective in mind: an organization wants its ‘uplink’ (the link between the organization’s network and its Internet Service Provider) to be transparent to the users, i.e., no negative impact on performance, see also, for instance, [6]. This objective will be achieved when the uplink’s bandwidth capacity C is chosen such that only during a small fraction of time ϵ , the aggregate rate of the offered traffic (measured on sufficiently small timescale T) exceeds the bandwidth capacity.

In more formal terms, this performance objective can be stated as follows. Let $A(T)$ denotes the (aggregated) amount of traffic offered in an (arbitrary) interval of length T . Then it is required that

$$\mathbb{P}(A(T) \geq CT) \leq \epsilon \quad (1)$$

For provisioning purposes, the crucial question is: “for given T and ϵ , what is the minimally required bandwidth $C(T, \epsilon)$ to meet the performance target?” In the remainder of this section we derive provisioning formulae that answer this question. First, we do this for general traffic (i.e., without modeling assumptions), then for the flow-level M/G/ ∞ model, and finally for (black-box) Gaussian models.

A. General traffic

Based on classical Markov inequality $\mathbb{P}(X \geq a) \leq (\mathbb{E}X)/a$ (for nonnegative random variables X), we find the following upper bound on the target probability – see (1) – by putting $X = e^{\theta A(T)}$, for $\theta \geq 0$:

$$\begin{aligned} \mathbb{P}(A(T) \geq CT) &= \mathbb{P}(e^{\theta A(T)} \geq e^{\theta CT}) \\ &\leq \mathbb{E}e^{\theta A(T) - \theta CT} . \end{aligned}$$

Because this upper bound holds for *all* nonnegative θ , we can choose the *tightest* upper bound, viz.:

$$\mathbb{P}(A(T) \geq CT) \leq \min_{\theta \geq 0} (\mathbb{E}e^{\theta A(T) - \theta CT}) . \quad (2)$$

This bound, also known as the *Chernoff bound*, is usually quite tight, but unfortunately rather implicit: it involves both the computation of the moment generating function $\mathbb{E}e^{\theta A(T)}$ as well as an optimization over θ .

To meet performance criterion (1), it suffices to choose C such that the right-hand side of (2) is below ϵ . It is easy to derive that the following *generic bandwidth provisioning formula* gives the lowest C to meet (1):

$$C(T, \epsilon) = \min_{\theta \geq 0} \frac{\log \mathbb{E}e^{\theta A(T)} - \log \epsilon}{\theta T} . \quad (3)$$

It is noted that in this generic bandwidth provisioning formula no modeling assumptions are imposed on the traffic, other than stationarity: as long as the distribution of the offered traffic at timescale T , i.e., $A(T)$, is known, the formula can be applied.

Essentially, formula (3) captures the following effects, which are intuitively clear: (i) the minimally required capacity consists of two parts, i.e., the average rate, plus some margin (following from Jensen’s inequality for convex functions); and (ii) C is decreasing in ϵ both as well as T (the looser the performance constraints, the less capacity is required).

In the sequel, we further develop the generic bandwidth provisioning formula, for two classes of models: M/G/ ∞ (‘flow-level’) and Gaussian (‘black-box’) input.

B. M/G/ ∞ traffic

In the M/G/ ∞ input model [8], flows arrive according to a Poisson process with rate λ , and stay in the system for some random duration D . During this ‘holding period’, traffic is generated at some rate. An obvious choice is to take this rate constant and deterministic (at some value $r > 0$). For instance, in a peak-rate limited environment (say, a ‘slow’ connection via a modem), r could be taken equal to the access rate: the ‘slow’ connection will always be the limiting factor for the rate at which traffic is generated, see for instance [5]. Alternatively, the rate could be constant but random: during each flow traffic arrives at a constant rate R , but this R is the realization of some random variable (and hence flows do not necessarily have the same rate). Finally, the M/G/ ∞ model even covers the case in which the rate at which traffic is generated during the ‘holding’ period follows some stochastic process, see e.g., [9], [10].

The M/G/ ∞ model is inherently flexible and comprehensive: by choosing specific flow-length distributions D , both short-range and long-range dependent inputs can be modeled (so-called heavy-tailed durations lead to long-range dependence). An important advantage of M/G/ ∞ modeling, is that it allows for ‘sensitivity’ analysis. For instance, in the model with constant and deterministic traffic rates, it is possible to predict the impact of changing traffic rates r of individual flows on the performance, which may be of particular interest in the area of bandwidth provisioning: “what is the effect of an upgrade of the network access speed on the aggregate traffic?”

In order to derive the minimally required bandwidth for M/G/ ∞ input, we determine the log-moment generating function in the generic bandwidth provisioning formula (3); for ease focus on the case of constant and deterministic traffic rates r . Let the mean flow duration $\mathbb{E}D$ be denoted by δ , such that the mean input rate $\mu = \lambda\delta r$. We denote by $F_D(\cdot)$ the distribution function of D , and by $F_{D^r}(\cdot)$ the distribution function of the residual flow length. The corresponding densities are denoted by $f_D(\cdot)$ and $f_{D^r}(\cdot)$. With $A(t)$ the amount of traffic generated by a single M/G/ ∞ input in an interval of length t , we distinguish between

- flows that were already active at the start of the interval. The number of these sources has a Poisson distribution with mean $\lambda\delta$. Their residual duration has density $f_{D^r}(\cdot)$; with probability $(1 - F_{D^r}(t))$ they generate traffic during the entire interval.
- flows that arrive during the interval. Their number has a Poisson(λt) distribution. Given that the number of these arrivals is a nonnegative integer, their arrival epochs are i.i.d. random variables, uniformly over the interval (with

density $1/t$). Their duration has density $f_D(\cdot)$.

Straightforward computations now yield the desired log-moment generating function, cf. [10]:

$$\begin{aligned} \log \mathbb{E} e^{\theta A(t)} &= \lambda \delta(M_t(r\theta) - 1) + \lambda t(N_t(r\theta) - 1), \quad \text{with} \\ M_t(r\theta) &:= \int_0^t e^{r\theta x} f_{D^r}(x) dx + e^{r\theta t}(1 - F_{D^r}(t)) \quad \text{and} \\ N_t(r\theta) &:= \int_0^t \int_u^t \frac{1}{t} e^{r\theta(x-u)} f_D(x-u) dx du + \\ &\quad \int_0^t \frac{1}{t} e^{r\theta(t-u)} (1 - F_D(t-u)) du. \end{aligned}$$

From the above we conclude that in order to determine the required bandwidth capacity, the model parameters r and λ , as well as the distribution D have to be known. There is a vast body of literature on the choice of these. Often r plays the role of the access rate, see for instance Ben Fredj *et al.* [5]; supported by extensive measurements D is often assumed to have a power-law tail, see for instance Crovella and Bestavros [2]. We also refer to the detailed study by Ben Azzouna *et al.* [7] and the references therein.

C. Gaussian traffic

We recall the notion of a ‘black-box’ model, in that it abstracts from modeling individual users or flows. A commonly used subclass of black-box models are the Gaussian models. Assuming that the traffic aggregate $A(T)$ contains contributions of many individual users, in many situations it is justified to assume that $A(T)$ is Gaussian if T is not too small, see e.g., [4], [3].

In other words, $A(T) \sim \text{Norm}(\mu T, v(T))$, where μ denotes the (long-term) average offered traffic rate. The log-moment generating function in the generic bandwidth provisioning formula (3) is then given by

$$\log \mathbb{E} e^{\theta A(T)} = \theta \mu T + \frac{1}{2} \theta^2 v(T).$$

Substitution in (3) and minimization over θ yields the following bandwidth provisioning formula for Gaussian traffic:

$$\begin{aligned} C(T, \epsilon) &= \mu + \min_{\theta \geq 0} \left(\frac{\frac{1}{2} \theta v(T)}{T} - \frac{\log \epsilon}{\theta T} \right) \\ &= \mu + \frac{1}{T} \sqrt{(-2 \log \epsilon) \cdot v(T)}. \end{aligned} \quad (4)$$

We note that (4) is in the same spirit as the equivalent bandwidth formula given in [11].

D. Test cases

In this study we make extensive use of ‘packet traces’ from operational IP networks. The traces are taken from five distinct networks, each with different traffic characteristics in terms of network access technologies (e.g., ADSL, Ethernet), link speeds (ranging from 512 kbit/s to 1 Gbit/s), number of subscribers, types of users (e.g., students, ‘normal consumers’), etc. These networks are selected to resemble various common real-life scenarios. Note that we have focused on traffic from and to the

access network (that is, to and from the core network (Internet), respectively), i.e., traffic that is sent over the ‘uplink’. LAN traffic as well as backbone traffic, has been subject of various other studies, see e.g., [1], [7].

The measurement procedure to gather the packet traces in our study is as follows. For each network, we have hooked up an off-the-shelf PC to a router/switch that copies all traffic from/to the network to the measurement PC. Using the standard tcpdump software, all packet headers are captured, time-stamped, and subsequently made anonymous through the tcpdpriv tool to protect the privacy of users; the procedure is detailed in [12]. In this way we have obtained over 400 traces in total, each of them containing 15 minutes of traffic. The traces are available online [13]. Because of space restrictions, we cannot discuss all traces or all networks here, and instead have to limit to two test cases. We stress that analysis of the other traces has not lead to results that are contradictory to the findings presented in this paper.

The first test case, referred to as *loc1*, considers the 1 Gbit/s uplink of a Dutch research institute. The institute employs about 200 researchers and support staff, who all have a 100 Mbit/s access link. The uplink is only mildly loaded – about 1% on average (long-term). The traces from this uplink have been gathered in Summer 2003. The second test case, referred to as *loc2*, considers the 1 Gbit/s uplink of an ADSL access network with several hundreds of subscribers. The ADSL access speeds vary from 512 kbit/s to 8 Mbit/s. The traces from this uplink were collected in Summer 2004. For the discussions in this paper, we have randomly selected traces from these locations.

In the next two sections we will use these test cases to investigate and illustrate the applicability of M/G/ ∞ (Section III) and Gaussian (Section IV) traffic modeling to support bandwidth provisioning.

III. M/G/ ∞ MODEL

As said before, a prominent advantage of M/G/ ∞ input modeling is that it allows sensitivity analysis of ‘user-level parameters’ – an attractive feature for bandwidth provisioning purposes. In this section we investigate the applicability of this model using the traffic traces described above, and we study to what extent it provides insight into the impact of the rate at which flows generate traffic. We first verify whether the basic version of the M/G/ ∞ model holds, i.e., the model in which r is constant and deterministic (i.e., all flows transmit at the same constant rate).

Heterogeneity of traffic flows may have various causes. In the first place, flows themselves may generate

Fig. 1. Accumulated traffic within four flows over time

traffic at fluctuating rates, various flows may have different (possibly constant) rates, etc. Differences in traffic rates among flows are the result of the use of various applications (that cause different traffic patterns), and the end-to-end nature of TCP/IP connections, in that any link within the path between the two end-points may reduce the achieved traffic rate.

In this section, we first investigate whether traffic *within a single flow* is generated at a (more or less) constant rate (Section II.B). Secondly, we compare traffic rates *among various flows* (Section III.B). Motivated by measurements that show that most of the traffic is generated by only a limited number of large flows ('elephants'), we then repeat this exercise only taking into account these long flows. As access rate limitations are imposed on the user's transmission rate rather than the transmission rate of an individual flow, we also aggregate all (concurrent) flows as generated by the same user and investigate whether these aggregated traffic flows streams exhibit homogeneity. Based on these observations we argue that the model with a constant and deterministic rate does not apply, and therefore we shift our attention to the model with a constant but random transmission rate (Section III.C).

A. Rate within a flow

Various definitions of a *flow* of network traffic are in use. At this point, we define a flow following the common 5-tuple definition: a flow comprises all packets with the same (*source IP address, destination IP address, transport protocol, source port, destination port*), as long as the 'gap' between such packets does not exceed some predefined interval. For instance, in this definition all IP packets within the same TCP-connection belong to the a single flow.

Figure 1 shows four (relatively large) flows, picked from the approximately 60,000 flows in the *loc1* trace. The slopes of the lines indicate the rate of the flows; the traffic rate of the 'fastest' flow is about eight times as high as the rate of the 'slowest'. The traffic rate within the flows, however, appears more or less constant (given the fairly straight lines). The differences in traffic rates may be explained by heterogeneity, for instance because the various flows stem from different applications with

Fig. 3. Cumulative distribution of flows and their contribution to all traffic

varying demands from the network. Another possible explanation for the different traffic rates may be that various end-to-end paths may have different bottlenecks that restrict the transmission rate.

This observation shows that one may not assume the traffic rate r to be a fixed value in $M/G/\infty$ modeling. The next subsection further study this heterogeneity.

B. Rate between flows

In this subsection we compare traffic rates between flows. We assume, motivated by the discussion above, that traffic rates within single flows are constant – every flow may have a different traffic rate though. To investigate the traffic rates between flows, we plot the duration of every flow against its size – their ratio is the traffic rate. Figure 2 shows all flows (according to the 5-tuple definition) in the *loc1* trace. Clearly, various flows of the same size, may take longer or shorter to complete. Similarly, the duration of a flow, does not provide us with any information on the flow's size; we observe an extreme heterogeneity.

One may wonder whether this heterogeneity may be caused by the above-mentioned mice-elephants dichotomy, or by the fact that we should aggregate flows per user. We now investigate these options.

The widely used assumption that Internet traffic is heavy tailed, motivates our choice to 'zoom in' larger flows. Therefore we first investigate which percentage of the flows cause what fraction of the traffic. Figure 3 clearly shows that only a small percentage of all flows account for most of the traffic. We decided to 'zoom in' on the approximately 3,000 largest flows (corresponding to 95% of all traffic, and about 5% of all flows). For this sub-set of all flows, the duration-size pair of each flow is plotted in Fig. 4; again, the spread of the (*duration, size*)-tuples suggests great heterogeneity.

A single user may have multiple flows generating traffic concurrently, e.g., he may be browsing the web while a file download is going on in the background. These flows may interact with each other with regard to the rate at which each flow generates traffic when (partly) following the same Internet path; in any case,

they share the access line. Such interaction may affect the ‘homogeneity’ in terms of rates of individual flows. Therefore, we aggregate flows generated by the same user that are ‘overlapping in time’, and with this new definition of flow we again plot the duration-size tuples. Note that we left out the 95% smallest flows, like above, because of their negligibly small contribution to the total traffic. With the new definition of flow, some 160 ‘aggregated flows’ remain. Figure 5 shows the resulting $(duration, size)$ -tuples. The cloud in Fig. 5 is not as dense as before, however, but the ‘spread’ is still considerable.

From the discussions above we conclude that traffic rates in trace *loc1* are not fixed between flows; this conclusion remains valid when only ‘elephants’ are considered, and when flows are aggregated per user. This conclusion is also supported by analysis of an extensive number of other traces, taken from the same and other networks; although the specifics of the achieved rates, and spread of flow sizes and durations may differ, the clouds suggest strong rate heterogeneity.

C. Random rate

The previous section showed that one cannot assume that the traffic rate is constant and deterministic. Therefore, we now consider a second option: the transmission rates are constant within flows, but the value R of this rate is random. In this subsection, we try to ‘fit’ real traffic in an $M/G/\infty$ model with random transmission rate R .

For every single flow, the rate R is determined by its size and duration, according to $S = RD$, where S denotes the flow’s size and D its duration. We first investigate whether R and D are independent. A necessary condition for independence is

$$\frac{\mathbb{E}D \cdot \mathbb{E}R}{\mathbb{E}S} = 1. \quad (5)$$

We compute $\mathbb{E}D$, $\mathbb{E}R$, and $\mathbb{E}S$ for 10 different traces taken from the same uplink as trace *loc1*, and plot the resulting fraction (5) in Fig. 6. Fraction (5) is also plotted for the set of all flows that together constitute 95% of all the traffic in every trace, leaving the majority of flows (i.e., the smaller flows) out (‘top contributing flows’). Clearly, in only 2 of the 10 traces the resulting fraction comes close to 1. Therefore, from Fig. 6 we conclude that R and D are *not* independent in these traces of real traffic.

The dependence between the rate and duration implies that the traces considered do not fit in the framework of the $M/G/\infty$ input model.

Of course one could split the aggregate traffic stream into various smaller ‘sub-streams’, on the basis of size, but also application, etc. Such an approach was pursued in recent studies on flow-level modeling, e.g., [7]. There it was found that it is possible to describe real traffic

Fig. 5. Flow durations compared to flow sizes (5% largest flows, aggregated per source IP address)

Fig. 6. Ratios $\mathbb{E}D \cdot \mathbb{E}R / \mathbb{E}S$

using an $M/G/\infty$ model, but the accuracy of the fit is at the expense of the number of sub-streams, and the corresponding tuning parameters (for instance the threshold that distinguishes the mice from the elephants). We remark that a lot of effort is put into grouping flows that are similar, for instance elephants, together. The extra effort that is required to estimate the modeling parameters accurately, may be unattractive to network operators. Also, when the nature of the traffic changes, for instance because of new popular applications, the estimation of the parameters has to be redone, which may require significant effort.

An other important remark is that we have not succeeded in recognizing the access rate in our traces: the transmission rates (which where, as said before, constant during the flow’s holding time) are apparently limited by other bottlenecks than the access rate. As a consequence, it appeared infeasible to do sensitivity analysis of the required bandwidth as a function of the access rate.

IV. GAUSSIAN MODEL

In Section III we have seen that the strong heterogeneity of the traffic in our traces appears to be a key problem to flow-level modeling. In this section we abstract from flow-level modeling, and turn to a black-box model, i.e., describing a traffic *aggregate*. We focus on the case of Gaussian traffic (as introduced in Section II.C).

First, we investigate if a Gaussian model accurately describes the traffic aggregate in our traces for $T = 1$ second, and determine a quantitative measure for the ‘goodness-of-fit’ (similar to the procedure followed in [3]) (Section IV.A). Second, we investigate whether a Gaussian traffic model may also be used to describe the real traffic on other timescales, and estimate the variance curve $v(\cdot)$ (Section IV.B). The variance function is then, third, used to showcase the provisioning procedure cf. (4) (Section IV.C).

We emphasize that the procedures described in this section are all but new – see e.g., [3]. The main goal of this section is the comparison with the procedure highlighted in Section III. There it turned out to be quite cumbersome to find an $M/G/\infty$ model with a good fit, whereas the present section indicates that it is relatively easy to find appropriate black-box models.

Fig. 7. Q-Q plot - trace *loc1* compared to Gaussian

Fig. 8. Variance over multiple timescales

A. Model fitting

We investigate whether the traffic in our traces is accurately described by a Gaussian process: $A(T) \stackrel{?}{\sim} \text{Norm}(\mu T, v(T))$. Note that literature suggests that this may be true for T not too small [3], [4]. We choose $T = 1$ second to start with, motivated by our expectation that timescales of this order are relevant for performance as perceived by end-users of interactive applications like web-browsing. In the next subsection, we also consider other timescales.

The estimates $\hat{\mu}$ and $\hat{v}(T)$ of the average and (sample) variance of the traffic rates in our traces can straightforwardly determined: $\hat{\mu} = 1/n \sum_{i=1}^n A_i$, and $\hat{v}(T) = 1/(n-1) \sum_{i=1}^n (A_i - \hat{\mu})^2$, where A_i denotes the amount of traffic offered in an interval of length T , and n the number of slots (i.e., 900 slots of 1 second). We note that the convergence of the estimator of the sample variance could be rather slow when traffic is long-range dependent [14, Ch. I]. We find that, for the *loc1* trace, $\hat{\mu} = 18.9$ Mbit/s and $\hat{v}(1) = 24.3$ Mbit²/s². With the regard to the *loc2* trace, these are $\hat{\mu} = 103.0$ Mbit/s and $\hat{v}(1) = 37.02$ Mbit²/s².

We use a so-called quantile-quantile plot (Q-Q plot) for testing the Gaussianity of the traffic. In this Q-Q plot, the order statistics are plotted against the inverse of the normal cumulative distribution function with mean $\hat{\mu}$ and variance $\hat{v}(T)$. The closer the point-pairs are to the diagonal in such a Q-Q plot, the more Gaussian the distribution of $A(T)$ is.

Figure 7 shows the comparison between the traffic in trace *loc1* and the Gaussian traffic model. Visually, the traffic seems to be ‘fairly Gaussian’, as most point-pairs are close to the diagonal. Note, however, how the Gaussian model fails to capture the head and tail of the distribution of $A(T)$. This motivates conservative bandwidth provisioning: if a Gaussian traffic model is assumed, one should be aware that the model is not accurate when traffic rates are relatively high.

In order to get a quantitative measure of *goodness-of-fit*, we use the *linear correlation coefficient* as defined in [3], which we denote with γ . Note that $-1 \leq \gamma \leq 1$, and $\gamma = 1$ means that the empirical distribution is identical to the model distribution. We find that $\gamma =$

0.994 for the *loc1* trace, supporting the ‘fairly Gaussian’ characterization above. It turns out that the *loc2* trace is similarly well approximated by Gaussian: $\gamma = 0.992$ (for time scale $T = 1$ second).

B. Variance estimation

In the previous paragraphs we have seen that, for $T = 1$ second, the traffic from our traces is well approximated by a Gaussian traffic model. We now look into other timescales, and also determine the variance $v(T)$ for T ranging from 0.01 to 128 seconds.

First, to investigate to what extent the *loc1* trace is Gaussian on various timescales, we compare the distribution of $A(T)$ in the *loc1* trace with $\text{Norm}(\mu T, v(T))$ by computing the linear relation coefficient γ for each T . The resulting goodness-of-fit is plotted in Fig. 8 (right axis). In line with other measurement studies (e.g., [4]), the *loc1* appears Gaussian for $T > 0.5$ seconds, and not-perfectly-Gaussian for smaller T . Also, the Gaussianity of the *loc1* traces reduces for $T > 10$ seconds, which may be partly caused by the relative low number of possible observations from the *loc1* trace for relatively large T .

Second, we compute the sample variance $\hat{v}(T)$ of trace *loc1*, as a function of T , with T ranging from 0.01 to 128 seconds. The results are also plotted in Fig. 8 (left axis). We compare these results with one of the key models in traffic modeling: fractional Brownian motion (fBm), which is a Gaussian model; see, e.g., [15] for more information on fBm. The variance function for fBm is given by $v(t) = \sigma^2 t^{2H}$, where H is the so-called Hurst parameter. For $H > 1/2$ this corresponds to long-range dependent traffic. We now fit (using the least-squares method) the sample variances from the *loc1* trace with the variance curve of fBm traffic; we find that with $H = 0.82$ and scaling with $\sigma = 5.38$, the sample and fBm variances are close to each other, for all inspected T .

C. Provisioning

We return to our primary objective of bandwidth provisioning (1). The analysis in the previous paragraphs has shown that it is justified to use a Gaussian model to describe the network traffic in our traces. Hence, the bandwidth provisioning formula (4) for Gaussian traffic can be used to determine the minimally required bandwidth capacity C to meet the performance criterion (1) for ‘link transparency’. In this subsection we study the accuracy of the resulting provisioning procedure.

We specify our performance criterion as ‘only during at most 1% of time, the rate per second of the offered traffic may exceed the available capacity’, i.e., $T = 1$ second, and $\epsilon = 0.01$. Using our estimates $\hat{\mu}$ and $\hat{v}(T)$, we estimate using (4) that the minimally required capacity to meet that criterion, for the *loc1* trace, is $C \approx 34$ Mbit/s. For the *loc2* trace, for which we specify $T = 1$ and

Fig. 9. Average traffic rate (per second) and required capacity for trace *loc1*

Fig. 10. Average traffic rate (per second) and required capacity for trace *loc2*

$\epsilon = 0.05$, the required capacity is estimated at $C \approx 118$ Mbit/s.

In order to assess the accuracy of the estimations for the required bandwidth capacity, we compare these figures with the actual traffic rates, averaged per second, that we derive from the packet header traces. The assessment is graphically shown in Fig. 9 and Fig. 10. Clearly, the computed required capacity is sufficient to meet the offered traffic in almost all intervals. One might argue that the estimated required capacity may be exceeded more often without violating the performance criterion (actually 9 intervals, instead of the about 3 intervals as it shows from Fig. 9) – which is true. A reason for the slight ‘overshooting’ may be in (2), in that the Chernoff bound which gives not an exact value but rather an upper bound on the overflow probability. This actually leads to a somewhat conservative capacity estimation – which we believe is to be preferred over a scarce capacity figure, because of its inherent effect on (user perceived) performance.

V. CONCLUDING REMARKS

We have presented a comparison between two fundamentally different approaches of traffic modeling: (i) ‘flow-level’ modeling, capturing the behavior of individual flows, and (ii) ‘black-box’ modeling, which statistically describe the superposition of many users. In particular, we compare the ‘flow-level’ M/G/ ∞ model and a Gaussian model as ‘black-box’. The application that we have in mind in this study is network link provisioning; in particular, we are interested in achieving link transparency, i.e., the link should not have a negative impact on the performance. An accurate traffic model helps, in that it facilitates the prediction of the offered performance as function of the link utilization. We provide link provisioning formulae for both traffic modeling approaches.

Our comparison between the two modeling approaches is supported by packet header traces that were obtained from various real-life settings. The analysis of these traces has shown that the strong heterogeneity often makes the use of ‘flow-level’ models (of the M/G/ ∞ -type) unattractive. Contrarily, Gaussian models fit nicely,

especially for timescales larger than 0.5 second. Subsequent validation of the provisioning formula for Gaussian traffic, using the traces, shows that an (somewhat conservative, but still) accurate estimation of the required capacity is obtained.

In future work we intend to further investigate the applicability of a Gaussian traffic model in real-life settings. Special focus will be on its limitations: up to which timescales and user aggregation levels can the model be used? Also, the dominant use of Internet changes over time, and this may have consequences to traffic modeling – whether or not a Gaussian traffic model is still appropriate is a topic that requires further research.

REFERENCES

- [1] W.E. Leland, M.S. Taqqu, W. Willinger, and D.V. Wilson, “On the self-similar nature of Ethernet traffic,” *IEEE/ACM Trans. on Networking*, vol. 1, no. 2, pp. 1–15, February 1994.
- [2] Mark E. Crovella and Azer Bestavros, “Self-similarity in World Wide Web traffic: evidence and possible causes,” *IEEE /ACM Transactions on Networking*, vol. 5, no. 6, pp. 835–846, 1997.
- [3] Jorma Kilpi and Ilkka Norros, “Testing the Gaussian approximation of aggregate traffic,” in *Proceedings of the 2nd ACM SIGCOMM Internet Measurement Workshop*, pp. 49–61, 2002.
- [4] Chuck Fraleigh, Fouad Tobagi, and Christophe Diot, “Provisioning IP Backbone Networks to Support Latency Sensitive Traffic,” in *Proceedings of IEEE Infocom*, April 2003.
- [5] S. Ben Fredj, T. Bonald, A. Proutiere, G. Régnié, and J.W. Roberts, “Statistical Bandwidth Sharing: A Study of Congestion at Flow Level,” *ACM SIGCOMM Computer Communication Review*, vol. 31, no. 4, pp. 111–122, October 2001.
- [6] Chadi Barakat, Patrick Thiran, Gianluca Iannaccone, Christophe Diot, and Philippe Owezarski, “Modeling Internet backbone traffic at the flow level,” *IEEE Transactions on Signal Processing*, vol. 51, no. 8, August 2003.
- [7] N. Ben Azzouna, F. Clerot, C. Fricker, and F. Guillemin, “Modeling ADSL traffic on an IP backbone link,” *Annals of Telecommunications*, December 2004, to appear.
- [8] M. Parulekar and A. M. Makowski, “M/G/ ∞ input processes: A versatile class of models for network traffic,” in *Proceedings of IEEE INFOCOM*, April 1997, pp. 419–426.
- [9] N. Duffield, “Queueing at large resources driven by long-tailed M/G/ ∞ -modulated processes,” *Queueing Systems*, vol. 28, pp. 245–266, 1998.
- [10] Michel Mandjes and Miranda van Uitert, “Sample-path large deviations for tandem and priority queues with Gaussian inputs,” Available at <http://www.cwi.nl/ftp/CWIreports/PNA/PNA-R0221A.pdf>, June 2004.
- [11] R. Guérin, H. Ahmadi, and M. Naghsineh, “Equivalent capacity and its application to bandwidth allocation in high-speed networks,” *IEEE Journal on Selected Areas in Communications*, vol. 9, no. 7, pp. 968–981, September 1991.
- [12] Remco van de Meent, Aiko Pras, Michel Mandjes, Hans van den Berg, and Lambert Nieuwenhuis, “Traffic Measurements for Link Dimensioning: A Case Study,” in *Proceedings of the 14th IFIP/IEEE Workshop on Distributed Systems: Operations and Management (DSOM2003)*, pp. 106–117, October 2003.
- [13] “M2C Measurement Data Repository,” <http://m2c-a.cs.utwente.nl/repository/>.
- [14] Jan Beran, “Statistical methods for data with long-range dependence,” pp. 404–416, Nov. 1992.
- [15] Ilkka Norros, “On the Use of fractional Brownian motion in the Theory of Connectionless Networks,” *IEEE Journal of Selected Areas in Communications*, vol. 13, no. 6, pp. 953–962, 1995.