Pragmatic Modeling of Broadband Access Traffic

Matthew Roughan and Charles Kalmanek
AT&T Labs – Research,
180 Park Av., Florham Park, NJ, 07932
{roughan,crk}@research.att.com.

Abstract—**Good traffic modeling is a basic requirement for accurate capacity planning. The recent discovery of heavy-tails, and Long-Range Dependence (LRD) in traffic has heralded a new, and more elegant way to model data traffic, particularly characteristics such as extreme burstiness across many time scales. However, most of the measurements used to populate such models have been fine grained packet traces. In reality we are far from being able to obtain such traces from more than a small subset of the Internet, and this is likely to remain true at least in the immediate future. The only source of ubiquitous data is SNMP, but SNMP has many limitations which make it difficult to work with for traffic modeling. These limitations make it impossible to use standard LRD models. However, we show here that for broadband access, SNMP is capable of capturing the most important features of the data traffic. We base this analysis on a large volume (more than two months) of SNMP data obtained from a large operating broadband access network. The model is approximate, but is nonetheless quite useful for capacity planning. The results validate our intuition about LRD in data traffic, while allowing the key parameters of the model to be computed solely from SNMP traffic utilization data.**

I. Introduction

Over recent years the discovery of self-similarity in data traffic (see [1], [2], [3]) has stimulated a great deal of research into performance models for self-similar, and heavy-tailed traffic (for instance see [4], [5]). However, the great majority of this work relies on sophisticated, and detailed models, derived from packet traces. Unfortunately, in practice, packet trace data is rarely available because collecting such traces requires dedicated equipment that would be very expensive to deploy universally. Furthermore, the cost of collecting, and storing such large volumes of data is itself prohibitive[1].

However, good traffic models are required for ongoing tasks such as capacity planning, in particular in access networks where it is prohibitively expensive to "just add more bandwidth" in order to fix performance problems. Furthermore, network operators also need to be able to determine whether a link is overloaded or not. In fact, even this simple task, which we shall refer to here as capacity management, is not as easy as it might at first appear. The fractal nature of data traffic makes such estimation inherently more difficult.

The only type of traffic data that is nearly universally available for tasks such as capacity planning is SNMP (Simple Network Management Protocol) data. Flow level, or even packet trace data is often available in a limited way (at a few sample points, or for limited time periods), and can therefore supplement SNMP data, but cannot replace it. SNMP data is generally limited to coarse sampling periods, and suffers from missing and noisy data as a result of the protocol design. Hence, any traffic model must be able to deal with such data quality problems.

The noisy/unreliable nature of this data mean that any approach we adopt will be a crude approximation, and capacity planning based on this data will likewise involve approximation. However, the current rapid growth and uncertainty in the Internet mean that this is all that is currently needed. In the future we can aim to improve the SNMP mechanisms used to collect traffic data so that when it becomes important to squeeze that extra 5-10% out of our networks we have the data to support it, but for the moment we shall be satisfied with a model that gives us approximate answers, using the available data.

This paper analyzes a large volume (more than two months) of SNMP data obtained from a large operating broadband access network. The data we have is derived from cable access, but there is no reason to believe that cable users generate traffic that is intrinsically different from DSL users. Cable and DSL network use is growing dramatically, and they provide an order of magnitude (or more) increase in bandwidth to the consumer than dial-up access. As a result, these networks require significantly more bandwidth in the metro and regional networks where traffic is aggregated and backhauled to the backbone network point of presence. The metro and regional networks supporting cable and DSL are too large to allow for gross over-dimensioning to be economically viable. Furthermore, the investment climate that contributed to Internet growth until late in 2000 no longer exists, and network growth including capacity upgrades are likely to be capital constrained for the foreseeable future. This leads to a need for simple traffic models that can be derived from available traffic measurements on these high-speed access networks. However, traffic on these networks is dramatically different from traffic on dial-up access networks and little work has been done so far to characterize high-speed access traffic.

We examine the SNMP cable traffic data, and derive traffic models from this data that can be used to perform approximate capacity management and planning (detailed capacity planning requires finer grained time series than SNMP gives us access to). The traffic models we present are consistent with the more detailed studies of packet traces which discovered and explained self-similarity and Long-Range Dependence (LRD) in data traf-

[1] It is in fact possible to obtain measurements of LRD and other quantities on-line, as data is collected [6], but this has not been widely used yet

fic. However, we validate the traffic models using SNMP measurements, and suggest that SNMP measurement could be used in operational networks to derive the salient parameters of the both individual user and aggregate cable traffic distributions. The major insight gained from this study is the significant impact that high-speed access has on the distribution of the carried traffic, as compared with typical dial-up traffic[2]

In the following section we shall give a brief overview of the relevant background to this paper, including what we mean by self-similar and LRD traffic, and the features of SNMP data which limit how it can be used. Section III describes the datasets and analysis resulting from this study. The main result is that a simple log-normal distribution can be used to represent the minute-to-minute variations in the traffic around a (relatively) slowly changing mean rate (changing on the scale of hours). Section IV concludes the paper with a summary of the findings.

II. Background

We begin this paper with the relevant background, in particular on self-similar traffic modeling, and SNMP traffic data.

A. Self-similarity and heavy-tails

Self-similarity and Long-Range Dependence [1] (LRD) provide a natural, quantitative, and yet wonderfully elegant way of describing and modeling data traffic. The simplest self-similar model, Fractional Brownian Motion (FBM) has only three parameters, and yet can describe some sets of traffic (including its burstiness) over scales ranging from milliseconds up to hours.

Self-similarity and long-range dependence are closely related phenomena, typically estimated using the *Hurst parameter*, H. Self-similarity relates to the fact that under suitable scaling, the statistics of the traffic are the same at any time scale. Long-range dependence rates to correlations in the data which, though decreasing over wider ranges, never become insignificant. We shall not dwell on the technical aspects of these properties in this paper, as LRD cannot be measured using the SNMP data we have available (see below). Moreover, the results in this paper indicate that a more detailed LRD model is not needed to explain the major characteristics of the traffic.

However, an understanding of the origins of LRD in data traffic is very informative here. We shall in fact find that the origins of the traffic characteristics we see in cable access traffic are one and the same as the origins of LRD in other traffic sets. Thus, although our measurements suggest an alternative model be used for cable access traffic, the findings of this paper are consistent with those in the prior traffic modeling literature.

The origins of LRD lie in heavy-tailed distributions. Such distributions have large weights in the tail, and may exhibit properties such as infinite variance, or infinite mean. An example is the Pareto distribution which is in a class of distributions with a power-law tail. That is, the distribution functions tail looks like $1 - F(t) \sim Lt^{-\gamma}$. Of particular interest here will be the case where $1 < \gamma < 2$, where the mean of the distribution is finite, but the variance is infinite.

[2] Although we do not directly compare our measurements here with dial-up traffic measurements, the intrinsic modem speed limit of dial-up traffic makes it impossible for dial-up traffic to have the characteristics we observe here.

A.1 ON-OFF models

One suggested model for the origin of LRD is the superposition of ON/OFF sources with heavy-tailed ON or OFF periods [7], [8], [9]. More formally, we can model a single source as a renewal process alternating between two states: “ON” and “OFF”, where the generated traffic rates are R and 0, respectively. Moreover, the duration of at least one of the ON and/or OFF periods has a heavy-tailed distribution with infinite variance, for example a power-law distribution with $1 < \gamma < 2$.

ON/OFF processes with heavy-tailed ON periods are asymptotically LRD in themselves [10] with $H = (3 - \gamma)/2$, where the heavy-tailed distribution has exponent γ. However, a better aggregate model may be obtained by superposing a number of sources. The exact way in which the processes are superposed does matter, principally through the renormalisation [11]. Renormalisation is used so that as the number of sources is increased, the average rate remains constant. For clarity, we simplify the more general description of [11]. The simplest renormalisations are:

(i) reduce the rate R of each source,
(ii) increase the length of OFF periods or
(iii) reduce the length of ON periods.

As the number of sources goes to infinity, each method results in a different model,

(i) Fractional Brownian Motion with $H = (3 - \gamma)/2$,
(ii) M/G/∞ source model, and
(iii) α-stable process,

respectively.

These different renormalisations are of particular interest here because the standard assumption has been that in aggregated traffic either the first or second case will be the suitable model. For instance, for dial-up traffic, which is tightly constrained by the bandwidth of the modem connection, data transfers are spread over time. We can easily see these data transfers as an On/Off process with a low rate per source, and previous measurements of such traffic have supported the FBM type of model. Another rationale for the FBM model is that on a heavily loaded link the TCP congestion control causes the available bandwidth to be shared (approximately) between sources in such a way that a file transfer will be extended over time.

Our results show that with moderately aggregated traffic on a high-speed access network, the third model is more appropriate. The results show that the marginal distributions of the traffic are quite inconsistent with either of the first two models, but are consistent with the third. Bandwidth limitations and TCP still slow down the effective transfer rate, “spreading” the connections over time, however the greater access bandwidth in a cable network reduces the amount of spreading so that the majority of data transfers take less time than the measurement intervals. We discuss the measurement intervals more in Section II-B.

A.2 α-stable distributions

There are a number of reasons why one cares about LRD and heavy-tails. The most common is performance. Systems carrying LRD traffic will typically have worse performance (see [4], [5] for example). However, although the main application of this work is capacity planning (for which performance is important), we must cope with the existence of LRD and heavy-tails even before we plug parameters into a capacity planning model, because the very presence of LRD and heavy-tails makes measurement of the model parameters much more difficult.

Heavy-tails and LRD disrupt the standard intuition about how data will be smoothed by averaging, either over time, or over additional sources. The standard intuition is based on the Central Limit Theorem (CLT), which says that as we sum more and more data (suitably renormalised) the result tends to a Gaussian distribution whose mean tells us the mean of our data. The problem is that the CLT has two conditions, that the data be independent, and have finite variance.

Most people assume that if data is averaged over longer and longer time intervals you will eventually reach a regime where the data correlations have dropped to the point where they are insignificant, and therefore we can apply the CLT. However if the data is LRD we know that we will never reach this regime, and so the CLT does not apply. Alternatively, if the data has heavy-tails with infinite variance, then once again the CLT does not apply.

There is a Generalized Central Limit Theorem (GCLT) that deals with these cases. In the case of sums of heavy-tailed data, the GCLT provides the following result [12].

Theorem II.1 (Generalized Central Limit Theorem) Let X_1, $X_2, X_3, \ldots$, be an independent, identically distributed series of random variables. There exist constants $a_n > 0$, and $b_n \in \mathbb{R}$ and a non-degenerate random variable Z, with

$$a_n(X_1 + X_2 + \cdots + X_n) - b_n \xrightarrow{d} Z,$$

if and only if Z is α-stable, in which case $a_n = n^{-1/\alpha}$ for some $\alpha \in (0, 2]$.

Thus the *α-stable* distributions are the generalization of the Gaussian distribution which allow for the sum of heavy-tailed distributions. Note carefully that the normalizing factor depends on α which gives the "heaviness" of the tails. When $\alpha = 2$ the stable distribution is the Normal distribution, and the theorem corresponds to the CLT, whilst when $\alpha < 2$ the tail is heavy, and the convergence happens more slowly. There is a similar theorem which can be applied in the LRD case, but we shall be mainly concerned with the version above.

The upshot of this result is that the confidence bounds for mean estimates converge more slowly than light-tailed random variables (when they converge at all), converging as $n^{-1/\alpha}$ instead of $n^{-1/2}$. Furthermore, the tails of α-stable distributions are heavier than those of Normal distributions, and so the confidence intervals for α-stable distributions will be wider in any case. In simple terms, as we average more and more data (say by averaging the rate over time) we won't see such simple convergence to the mean. The result will still appear bursty, even over quite long time scales.

Throughout this paper we shall use the S0 parameterization of the α-stable distribution defined by the following.

Definition: A random variable X is said to be α-stable with S0 parameterization $S(\alpha, \beta, \gamma, \delta; 0)$ if and only if

$$X = \begin{cases} \gamma\left(Z - \beta \tan\frac{\pi\alpha}{2}\right) + \delta, & \alpha \neq 1, \\ \gamma Z + \delta, & \alpha = 1, \end{cases}$$

where $Z(\alpha, \beta)$ is defined by its characteristic function

$$E\left[\exp(iuZ)\right] = \begin{cases} \exp\left(-|u|^\alpha\left[1 - i\beta\tan\frac{\pi\alpha}{2}(\operatorname{sign} u)\right]\right), & \alpha \neq 1, \\ \exp\left(-|u|\left[1 + i\beta\frac{2}{\pi}(\operatorname{sign} u)\ln|u|\right]\right), & \alpha = 1, \end{cases}$$

where $\alpha \in (0, 2]$ is the *characteristic exponent*, $\beta \in [-1, 1]$ is the skew parameter, $\gamma > 0$ is the scale parameter, and δ is the location parameter, and the sign function indicates the sign of the argument (and is zero if the argument is zero).

The characteristic function above precisely defines the distribution, but requires numerical methods to compute densities as closed form expressions for the densities do not exist (except in some special cases). See [13], [14] for more details.

B. Measurement and estimation

Even the best methods for estimating LRD parameters (such as the Abry-Veitch wavelet based methods [15]) require traffic data on sub-second intervals[3]. Although that particular method may be adapted to operate on-line, without storing huge volumes of data [6], this still requires some type of measurement box to be placed into the network to collect and analyze the data. Whilst this is practical in the small scale, it is impractical in large operational networks to instrument every link in this way.

There are some alternative methods used to make measurements of traffic. The chief amongst these are based on flow statistics (for instance see Cisco Netflow [16], or RTFM [17]) or SNMP [18], [19] measurements. The former are not intended for deriving time series, and the latter inherently produce time series on coarse time scales (1 minute and greater). Neither method is therefore suitable for measuring LRD.

The question then, is can anything be done, in practice, with the available coarse data. This paper concentrates on the use of SNMP data because this is most widely, and easily collected. We must first understand exactly what we are getting with this data.

[3] The Abry-Veitch estimator which has close to the minimum variance estimate of the parameter H (and many other advantages in terms of robustness, and computational complexity) still requires thousands of data points, if not 10s or 100s of thousands for reasonable estimates. Given that network data shows strong non-stationarity over a period of a day, a huge number of measurements are required to get a Hurst parameter estimate better than 0.5 to 1, the allowed range.

SNMP operates through polling – a network management station polls the network elements (such as routers). SNMP can gather much more than just traffic data (for instance information about link state), however we shall only concern ourselves with traffic data. The traffic data at the network element is stored in the form of counters

1. ifInOctets
2. ifOutOctets

which store the number of bytes input and output to an interface on the router, respectively. Note that these are total counts, and are not reset by polls, and so the traffic rate on the link is the derivative of the counter value. When the counter reaches its maximum value it simply wraps around (without giving any notification) and starts at zero again. However, this is not the only circumstance in which a counter will reset to zero. The counters may reset to zero under other circumstances as well, for instance if the router is rebooted.

There are many problems with SNMP data:

1. **inaccurate timestamps:** Delays in the network, timestamps from clocks which might not be synchronized, and delays in routers may lead to inaccurate timestamps, and hence to a low sampling rate.
2. **missing polls:** Polls are carried by UDP and may therefore be lost in the network, at the router, or at the poller.
3. **multiple counter wraps:** In SNMP v1 the counters are only 32 bits and may therefore wrap multiple times in one polling interval.
4. **implementation issues:** Some implementations of SNMP do not appear to behave in the standard manner, particularly on devices like cable modems where this aspect of the modem is often not considered important.

The result of these impairments is a coarsely sampled data set, with some noise, and variable amounts of missing data. A great deal of care needs to be taken to even get data of this quality, and it is easy to make mistakes in processing the data. The important upshot of these data limitations is that it is impossible to make reasonable estimates of parameters of LRD (such as the Hurst parameter) with SNMP measurements. However, given that SNMP link utilization data is widely available in IP networks, if it is possible to use it to gain meaningful insight into the dynamics of access network traffic, then the data has the potential to be a valuable tool for capacity planning and management.

III. SNMP MEASUREMENTS AND ANALYSIS

This section describes the SNMP utilization data used in this study. AT&T Broadband collected the SNMP data from six Cable Modem Termination Systems (CMTS), and a sample of approximately 20% of the individual Cable Modems (CM) on these CMTSs, during January 18th to March 20th 2000. The data was collected at approximately one minute intervals.

The most notable features of this type of cable access network are:

- high speed access (the system studied here used a 10 Mbps shared channel in both the upstream and downstream directions)
- low level of aggregation (200-300 homes per CMTS)
- high variability in demand among customers
- often asymmetric traffic pattern (not always)
- ”always on” service

One characteristic of Internet traffic is that new applications, such as peer-to-peer applications, can appear overnight and dramatically change traffic patterns. During the period when our data was collected, the dominant consumer Internet application was Web access. Thus, our results apply most strongly to this type of traffic. We also note that while we study cable traffic here, there is no fundamental difference between cable and DSL from a traffic perspective[4]

Figure 1 shows a simplified network architecture, to illustrate where the measurements are made. The CMTSs aggregate traffic from a number of homes passed by the hybrid fiber coax plant, and the traffic is then sent to an aggregation router, which takes traffic from several CMTSs. Table I shows the numbers of houses connected to each of the CMTSs used in this study. The CMTSs and cable modems have been polled at both Ethernet, and RF interfaces. This redundancy provides a check on the data integrity, but we shall only report data from the one interface of the CMTS and the CM. The traffic data at the (Ethernet side of the) CMTS gives the aggregate traffic from all of the CMs connected (minus some small local traffic), while the traffic data from each CM gives the traffic to that subscriber individually.

Each of the homes passed might be using the cable network to provide cable TV, telephone, high-speed data, or some other service. Table I shows how many of the total number of homes connected to the network were subscribers to the high-speed data service. Each subscriber has a CM, but only a sample of these were polled by the SNMP poller to reduce the overhead due to polling. Table I shows how many CMs were polled for each CMTS used in the study below.

A. CMTS data

Figures 2 and 3 show the downstream and upstream[5] traffic rates respectively. The data gives the total aggregate traffic from all of the CMs connected to that CMTS.

Note that this data was from early 2000 before the emergence of peer-to-peer applications - therefore the traffic is highly asymmetric. Of the six CMTSs polled there were significant problems with two data-sets (c, and d), namely large amounts of missing, and obviously corrupted data (see above for reasons), leading one to suspect the remaining data is also corrupted. The remaining four CMTSs (sets a, b, e, and f) form the data used here.

The first task is to examine the data for non-stationarity. A

[4] The argument that cable access is shared and DSL is not is spurious, since DSL traffic is aggregated at the access multiplexer in any event.

[5] We use downstream to denote traffic flowing out of the backbone, to the customer, and upstream to denote traffic flowing from the customer to the backbone.

cursory glance shows what appears to be a daily cycle in the data. We explore this by plotting the average traffic rate at each time of day (by hours) over the entire data set. Figure 4 shows time of day averages for the four data sets.

On an ISP backbone, traffic follows a daily pattern fairly closely. However, in the broadband traffic we observe that the variation between days is actually larger than the variation within a day (as shown in Figure 4). For instance see Figure 5 which shows the 1 hour average of the traffic during the busy hour of each day. This large variation requires special care to be taken in the analysis of the data.

We next examine the marginal distribution of the traffic on these four CMTSs. The marginal distribution is the probability distribution of the 1 minute traffic rates ignoring the fact that the data is a time series, and treating each data point as if it were independent of the others.

However, since the data is non-stationary, we must first remove the effect of the non-stationarity as it would otherwise corrupt any further analysis. There are a number of ways to normalize the data to remove non-stationarity, based on the model of the traffic which one adopts. For instance, one could use the models:

- **additive mean:** The traffic rate
$$X(t) = m(t) + T(t).$$
- **multiplicative mean:** The traffic rate
$$X(t) = m(t) \times T(t).$$

where m is the mean traffic rate, and $T(t)$ is a stationary stochastic process representing the variation in the traffic rates. In the first case the appropriate way to normalize the traffic is to subtract the mean, and in the second to divide by the mean. For the CMTS data examined the second approach produced traffic rate statistics whose histograms were consistently of the same simple shape, with very similar parameters (see later), whereas the first approach produced complex multi-mode distributions that were different for each case. We concluded from this that the second model is appropriate here. In practice we don't know the mean of the traffic, but we estimate this either by blocking the data into groups, or using a moving average. We found that using a 1 hour block or moving window size for the moving average both produce consistent results suggesting that over periods of 1 hour the data is reasonably stationary. Significantly longer block sizes lead to inconsistent results.

The probability densities of the normalized marginal distributions of the traffic are shown in Figure 6. The y-axis shows probability, and the x-axis shows the normalized rate – hence $x = 1$ is where the mean rate over 1 minute is equal to the 1 hour mean, $x = 2$ means the 1 minute rate is twice the 1 hour rate, and so on. The graphs also show a simple moment based log-normal fit to the data and a Maximum Likelihood Estimate (MLE) based α-stable fit to the data [13], [14] (note that moment based fits are quite accurate and easy to compute for the log-normal distribution, because in the log domain the data is normally distributed, and so in this domain we compute the mean and variance, but the equivalent for an α-stable is not as simple or effective).

It appears from these graphs that the log-normal distribution provides a reasonable first order fit to the data (particularly on the right hand side of the mode of the distributions which is in any case the most important part of the distribution). However, the α-stable fit is better. Table II gives the parameters of the α-stable fit. Notice there is a relatively small range to the parameters with the variation in α being most significant.

It is hard to visually assess the goodness of fit of the distributions using a density graph, and so we also provide in Figure 7 a set of pp-plots which illustrate the goodness of fit. A good fit in these graphs will lie along the dotted diagonal line, while a poor fit would lie off the line. Once again we can see that the log-normal fit is not unreasonable, but that the α-stable fit is better.

One conclusion that we can draw is that the aggregate CMTS traffic data can be closely (though not perfectly) modeled using an α-stable distribution, and that in the cases we examined this distribution is reasonably close to a log-normal distribution (although that is not true of the α-stable distribution in general). We suspect that the α-stable distribution arises from the sum of a series of heavy-tailed random variables. We investigate this possibility further in the following section using traffic data from individual CMs.

B. Cable Modem Data

Our data includes a subset of the CMs associated with each of the CMTSs, giving us a view of the individual user's activity. Although data was collected from CMs associated with all of the CMTSs, we focus on the data from CMs associated with the four CMTSs for which we have good data (data sets a, b, e and f). In particular the results displayed show data set (a), but are representative of the results for all four data sets (though the details differ). The number of cable modems polled for each is given in Table I.

There are a number of things that we could do with the CM data, including investigating the variation among subscribers. However, in this paper we focus on developing results that can be directly used in traffic modeling and capacity planning.

First of all let us look at the marginal distributions for each subscriber. Figure 8 (a) shows the marginal distribution for each subscriber (for data set a) – note that these marginals are not normalized. In all but two cases we note that the bodies of the marginals (between around 200 and 6000 bps) are the same. This part appears to be some type of background traffic load or "chatter" since it occurs almost all the time, and is the same for most of the modems. Figure 8 (b) shows the average of the individual marginals both during the busy hour and over all times of day. Notice that the selection of time of day has no significant effect on the form of the marginal distribution. This is another indication that the marginal is a useful indicator of behavior.

The variation in the marginal distribution comes through the highly variable tails of these distributions. Notice that Figure 8 is a log-log plot, and that the tail in plot (b) appears to be linear, and so follows a power-law. Hence we have a heavy-tailed marginal distribution for these sources. This confirms our earlier suspicion that the α-stable distribution in the marginal of the aggregate traffic arises from the sum of the heavy-tails in the individual user traffic.

A standard model of user activity is an ON/OFF model, with fixed ON rate and ON/OFF times given by heavy-tailed random variables. We now analyze the CM data in the context of this model. We will call a subscriber ON if it sends/receives some user data during the one minute interval between polls. Note, however, that in our current real-world setting, we need to think carefully about what is user data. For example, does user data include periodic polls (or a ticker) from a mail/news/web browser? Does it include network operations data such as NTP? Certainly it should exclude network maintenance data between modems and other network elements, such as our own SNMP polling. It is not possible to distinguish the source of the traffic with SNMP data so we use a simple rate threshold here. As a result, we are unable to detect cases when a user sends a very small amount of data. However, since we are primarily interested in the large volume traffic, missing low rate traffic will not have a significant effect on the overall analysis.

The choice of threshold used to decide if a subscriber is ON is obviously important. Figure 9 shows the effect on the ON probability: overall (dashed) and for the busy period (solid), of changing the threshold. Clearly the curve has a knee at just below 5 kbps. Below this a change in the threshold has a dramatic effect on the ON probability, but above this the ON probability is relatively insensitive to the value of the threshold. We use a threshold of 6 kbps below, but as noted, the exact value has little effect on the results.

Figure 10 shows the probability of being ON for each of the cable modems, averaged over all time (solid), and during the busy hour (dashed). We can see that the variation between customers is larger than the variation between busy hour and non-busy hour. The plot also shows the conditional probability that a customer is ON in a given hour given it is ON at least once.

We can also study the relationship between the ON probability and the data rates. Figure 11 shows the average ON probability by time of day over the whole data set, and the averaged data rate of all of the sampled cable modems. Clearly the two are highly correlated. The fact that the data rate appears to have had a constant added (above the ON probability) is the result of the chatter between modem and CMTS. The correlation leads one to believe that the ON probability is the fundamental parameter determining the underlying rate of the traffic.

Notice that Figure 11 does not agree with Figure 4 (a). In fact this is no surprise. Recall that our CM data includes only a subset of the CMs associated with a CMTS. When dealing with heavy-tailed distributions, there is a high probability that a random sampling will miss the small number of heavy-hitters and therefore see a skewed version of the results, as we do here.

Note that a more complicated user model would be to have a notion of "active" and "inactive" subscribers, where an active subscriber is one that is using a network application in some way, but are not necessarily downloading something at that instant. There are hence a number of states for a user (inactive, active (but OFF), and active and ON). The appropriate model for activity pattern of a user would depend heavily on demographic data (is the customer a business, an adult consumer, a teenager?). However, the model itself (average ON probability, download size distributions) can be separated from the demographics and could be measured through mechanisms such as Netflow at aggregation routers (though this might only be done on a sampled basis because of the difficulties in collecting flow level measurements). This model provides an explanation for the dotted curve in Figure 10 which is the conditional probability that a source is ON, given that it is ON some time in the hour. Clearly this probability is much higher than the overall ON probability leading to the activity based model above. While we are interested in this model, we do not explore it further in this paper.

C. Explanations

We can explain the above measurements in the following way. As noted above the standard model of user activity is the ON/OFF model. The measurements of individual cable modems supports this view, with one notable variation which is that the ON rate is now so highly variable that it is more important than the correlations over time. The result is, that when superposed the traffic appears to have an α-stable marginal distribution (which can be approximated to some degree by a log-normal distribution).

Although the individual ON rates of users have a high level of variability, the average ON probability is small (~ 0.05), and the number of customers large enough so that we can approximate the number of ON sources by a Poisson random variable. When we do this, only the average number of ON customers matters (given by the ON probability times the number of customers).

A simple simulation of this type of process (Poisson sum of heavy-tailed distributions) results in the marginal shown in Figure 12 which matches well with that seen at the CMTSs. Note that the results of the simulation displayed in Figure 12 are not perfect. The simulation uses the simple ON/OFF model, and so does not take into account:

- the fact that sources are not homogeneous,
- correlations in the data (for instance the duration of the ON and OFF periods also appears to have a heavy-tail here),

which may effect the results. Figure 13 shows a scatter plot of the time series for data set (a) versus a shifted version of itself. The plot shows a clear correlation between the traffic in consecutive intervals. If we wish to obtain a more accurate simulation model we must include a model for these correlations. A nice way to do so is to use a generalization [11] of the simple normalization schema described in Section II-A.1 which allows both heavy-tails and correlations, however, the current SNMP measurements do not support such sophisticated analysis with any ease. Fortunately, these correlations seem to be a second order effect.

The other complicating factor, the non-stationarity in the data can be taken into account by varying the ON probability. Figure 11 shows the clear correlation between this ON probability and the average rate. A simulation such as the one above may then be used to perform capacity planning at a coarse level, though detailed capacity planning requires estimates of fine time scale behavior that are not immediately available from SNMP data.

IV. Conclusion

In this paper, we analyze a large amount of SNMP utilization data from a set of cable modem termination systems and associated cable modems. Despite the limitations SNMP data, we develop a simple model for the aggregate traffic at a CMTS, showing that the minute-to-minute variation in the average rate of can be approximated by a log-normal distribution, under a suitable renormalisation. Our results show that renormalizing by dividing the aggregate rate by a one hour moving window estimate of the mean rate gives quite good results. In addition, we show that this model can be explained by a more accurate approximation to an α-stable distribution. This model can in turn be derived from a superposition of a set of individual subscribers described by a simple subscriber model, where each subscriber acts as an ON/OFF source that transmits at a highly variable rate when ON.

The most appealing thing about this model is that it is possible to directly measure it from coarse grained, but readily available SNMP utilization data. Moreover, the model is consistent with the explanations that have been given for LRD in other traffic studies. This link strengthens our conclusions, and at the same time provides great insight into the difference between high and low speed access traffic.

These models can also interesting because of their potential application to capacity management and planning in broadband access networks. For example, when analyzing a cable network, it may be natural to ask whether it is possible to add another 100 users to an existing CMTS. By collecting and analyzing SNMP utilization data, one can estimate the parameters associated with the individual user traffic distribution. One can then estimate the effect of adding another 100 subscribers by simulating the superposition of a system with additional subscribers. Note that in practice, one wouldn't need to measure individual cable modems. A more practical approach would use other sources of information such as a sample of Netflow measurements to examine the distribution of the size of customers' flows, and use this to estimate the parameters in the distributions used here.

We note that one might conceivably obtain even better fits to the data using a distribution such as the Weibull distribution. However, the aim of this paper is not distribution fitting, but providing simple, pragmatic models. Obtaining a perfect fit is less important than having a model whose basis can be derived from simple assumptions about the underlying traffic, making it possible to easily extend the model to cases where there are no measurements. For instance, to answer "what if" questions such as what if we build a new network from scratch, with N subscribers per CMTS.

Acknowledgments

We would like to acknowledge the help of Cheryl Persinger and Jennifer Regensburger (nee Fess) of AT&T Broadband in collecting the data used in this study. We would also like to thank John P. Nolan for the use of his stable-distribution parameter estimation code [13], [14], and Walter Willinger for many helpful comments.

References

[1] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Transactions on Networking*, vol. 2, pp. 1–15, Feb 1994.

[2] V. Paxson and S. Floyd, "Wide-area traffic: The failure of poisson modeling," *IEEE/ACM Transactions on Networking*, vol. 3, no. 3, pp. 226–244, 1994. Available at `http://www.aciri.org/floyd/papers.html`.

[3] W. Willinger, M. Taqqu, and A. Erramilli, in *Stochastic Networks: Theory and Applications*, eds. F. P. Kelly and S. Zachary and I. Ziedins, ch. A Bibliographical Guide to Self-Similar Traffic and Performance Modeling for Modern High-Speed Networks, pp. 339–366. Clarendon Press (Oxford University Press), Oxford, 1996.

[4] A. Erramilli, O. Narayan, and W. Willinger, "Experimental queueing analysis with long-range dependent packet traffic," *IEEE/ACM Transactions on Networking*, vol. 4, pp. 209–223, April 1996.

[5] I. Norros, "On the use of fractional Brownian motion in the theory of connectionless networks," *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 6, pp. 953–961, 1995.

[6] M. Roughan, D. Veitch, and P. Abry, "Real-time estimation of the parameters of long-range dependence (extended version)," *IEEE Transactions on Networking*, vol. 8, no. 4, pp. 467–478, 2001.

[7] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson, "Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level," *Proceedings of the ACM/SIGCOMM'95*, 1995. Available at `http://www.acm.org/sigcomm/sigcomm95/sigcpapers.html`.

[8] M. S. Taqqu, W. Willinger, and R. Sherman, "Proof of a fundamental result in self-similar traffic modeling," *Computer Communications Review*, vol. 27, pp. 5–23, 1997.

[9] W. Willinger, V. Paxson, and M. S. Taqqu, "Self-similarity and heavy tails: Structural modeling of network traffic," in *A Practical Guide to Heavy Tails: Statistical Techniques and Applications* (R. Adler, R. Feldman, and M. S. Taqqu, eds.), pp. 27–53, Birkhauser, Boston, 1998.

[10] B. K. Ryu and S. B. Lowen, "Point process approaches to the modelling and analysis of self-similar traffic - part I: Model construction," in *IEEE INFOCOM'96*, vol. 3, (San Francisco, California), pp. 1468–1475, March 1996. Available from `http://www.wins.hrl.com/people/ryu/papers/index.html`.

[11] J. Levy and M. S. Taqqu, "Renewal reward processes with heavy-tailed inter-renewal time and heavy-tailed rewards," *Bernoulli*, vol. 6, no. 1, pp. 23–44, 2000.

[12] J. Beran, *Statistics for Long-Memory Processes*. Chapman and Hall, New York, 1994.

[13] J. P. Nolan, "STABLE - MLE for stable distributions." Available at `http://academic2.american.edu/~jpnolan/stable/stable.html`.

[14] J. P. Nolan, "Maximum likelihood estimation of stable parameters," in *Levy Processes* (O. Barndorff-Nielsen, T. Mikosch, and S. Resnick, eds.), Birkhauser, Boston, 2001. Available at `http://academic2.american.edu/~jpnolan/stable/stable.html`.

[15] P. Abry and D. Veitch, "Wavelet analysis of long-range dependent traffic," *IEEE Trans. on Info. Theory*, vol. 44, no. 1, pp. 2–15, 1998.

[16] "Cisco netflow." `http://www.cisco.com/warp/public/732/Tech/netflow/`.

[17] N.Brownlee, "RTFM: Applicability statement." IEFT, Request for Comments: 2721, 1999.

[18] J.D.Case, M.Fedor, M.L.Schoffstall, and C.Davin, "Simple network management protocol (SNMP)." IEFT, Request for Comments: 1157, 1990.

[19] J.Case, K.McCloghrie, M.Rose, and S.Waldbusser, "Protocol operations for version 2 of the simple network management protocol (SNMPv2)." IEFT, Request for Comments: 1905, 1996.

Data set	homes passed	subscribers	CMs polled
a	1071	194	35
b	790	176	36
e	833	250	50
f	1051	309	60

TABLE I

NUMBER OF CABLE MODEMS POLLED AT EACH CMTS

Data set	N	S0 parameters			
		α	β	γ	δ
a	3159	1.50	1.00	0.27	0.78
b	3095	1.70	1.00	0.32	0.82
e	3166	1.48	0.87	0.24	0.82
f	3159	1.40	1.00	0.30	0.75

TABLE II

MLE OF THE PARAMETERS (USING THE S0 PARAMETERIZATION OF NOLAN [13], [14]) OF THE α-STABLE DISTRIBUTION. N REFERS TO THE NUMBER OF DATA POINTS USED IN EACH FIT.

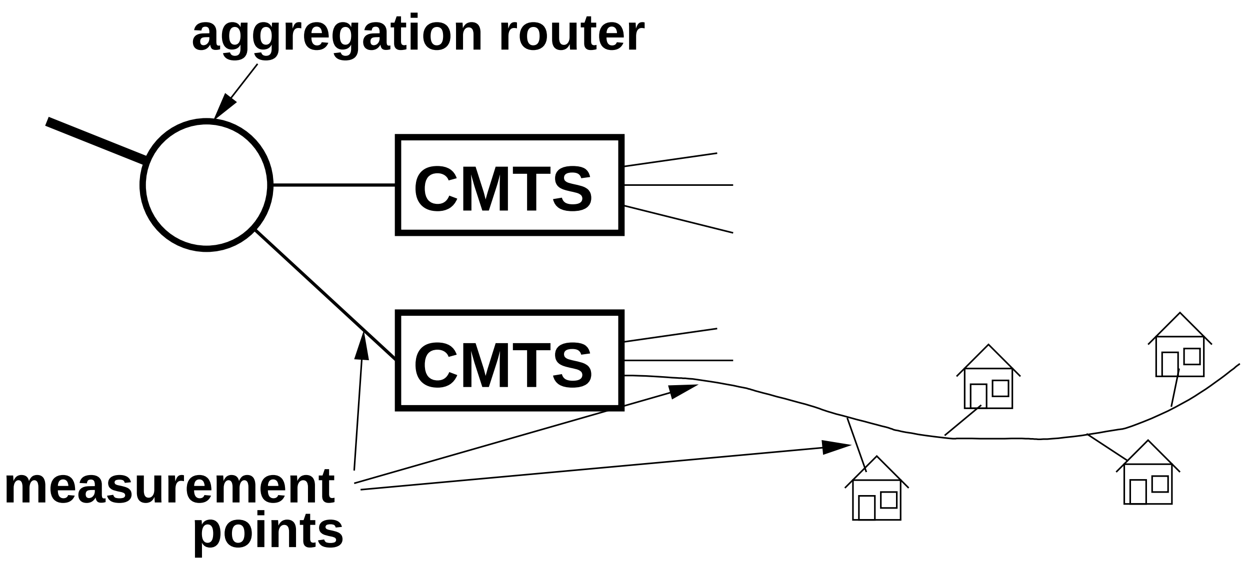

Fig. 1. Simplified architecture of the system in question. Our study uses data taken from the upstream (Ethernet) interface of the CMTS, the downstream (RF) interface of the CMTS, and from the CMs located at residences.

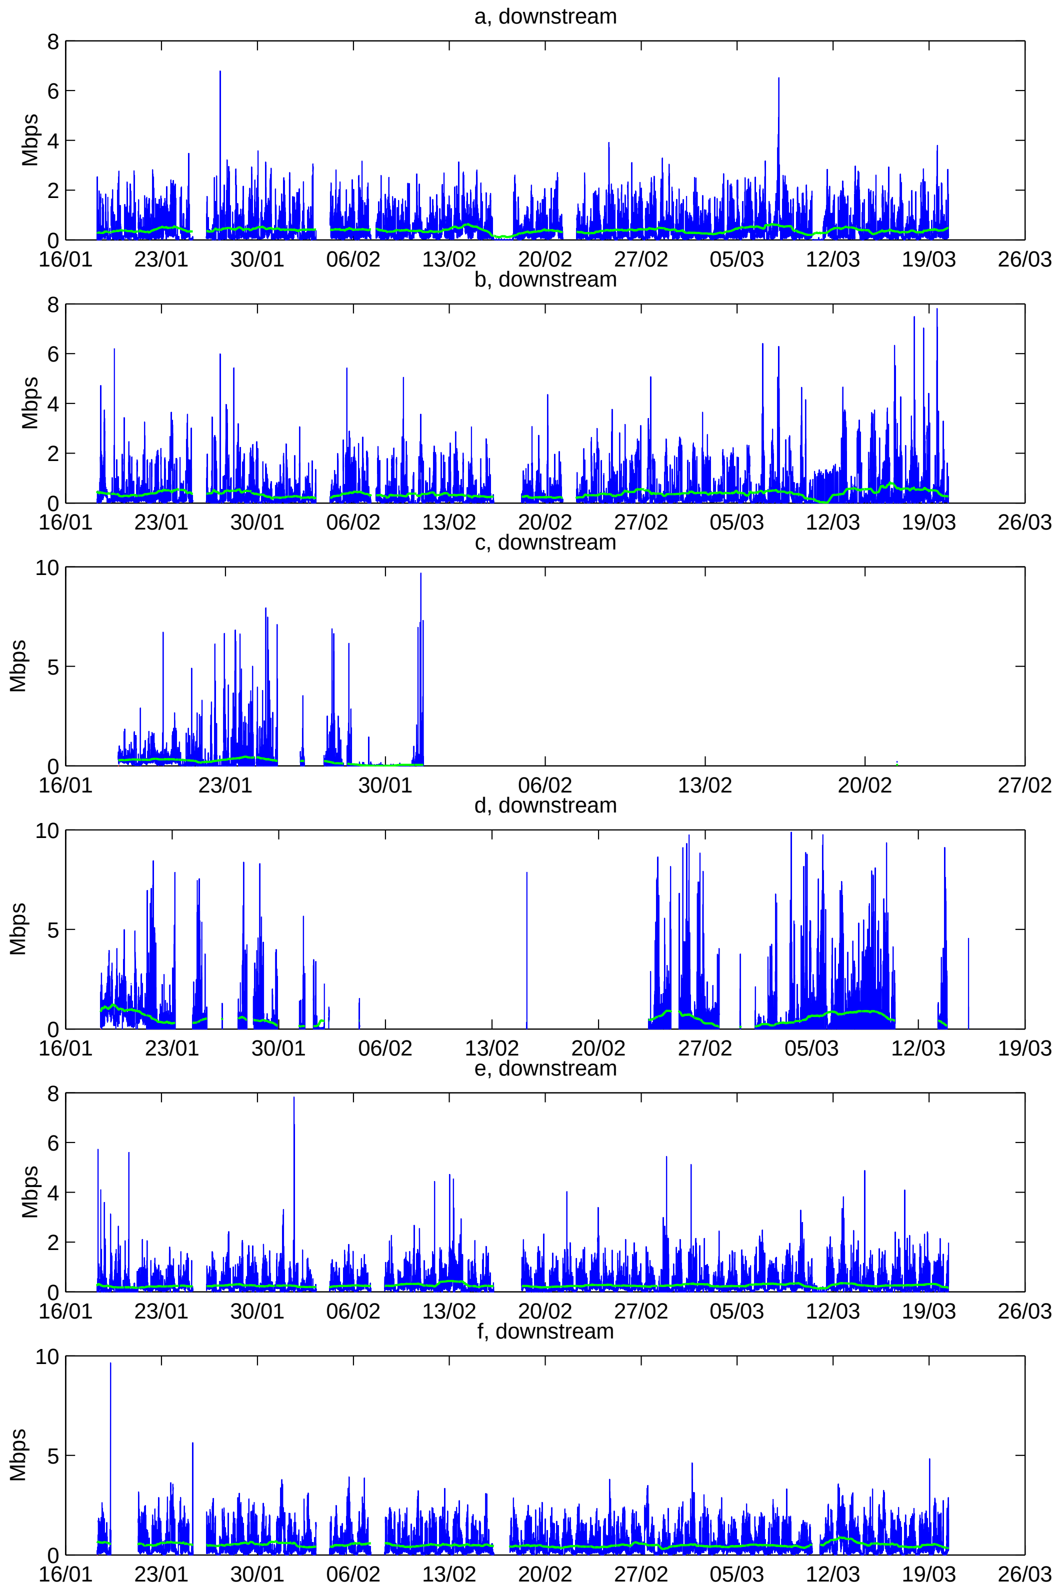

Fig. 2. Downstream CMTS traffic. The solid line shows one minute average rates (in Mbps) and the pale line shows a two week moving average of the data rate.

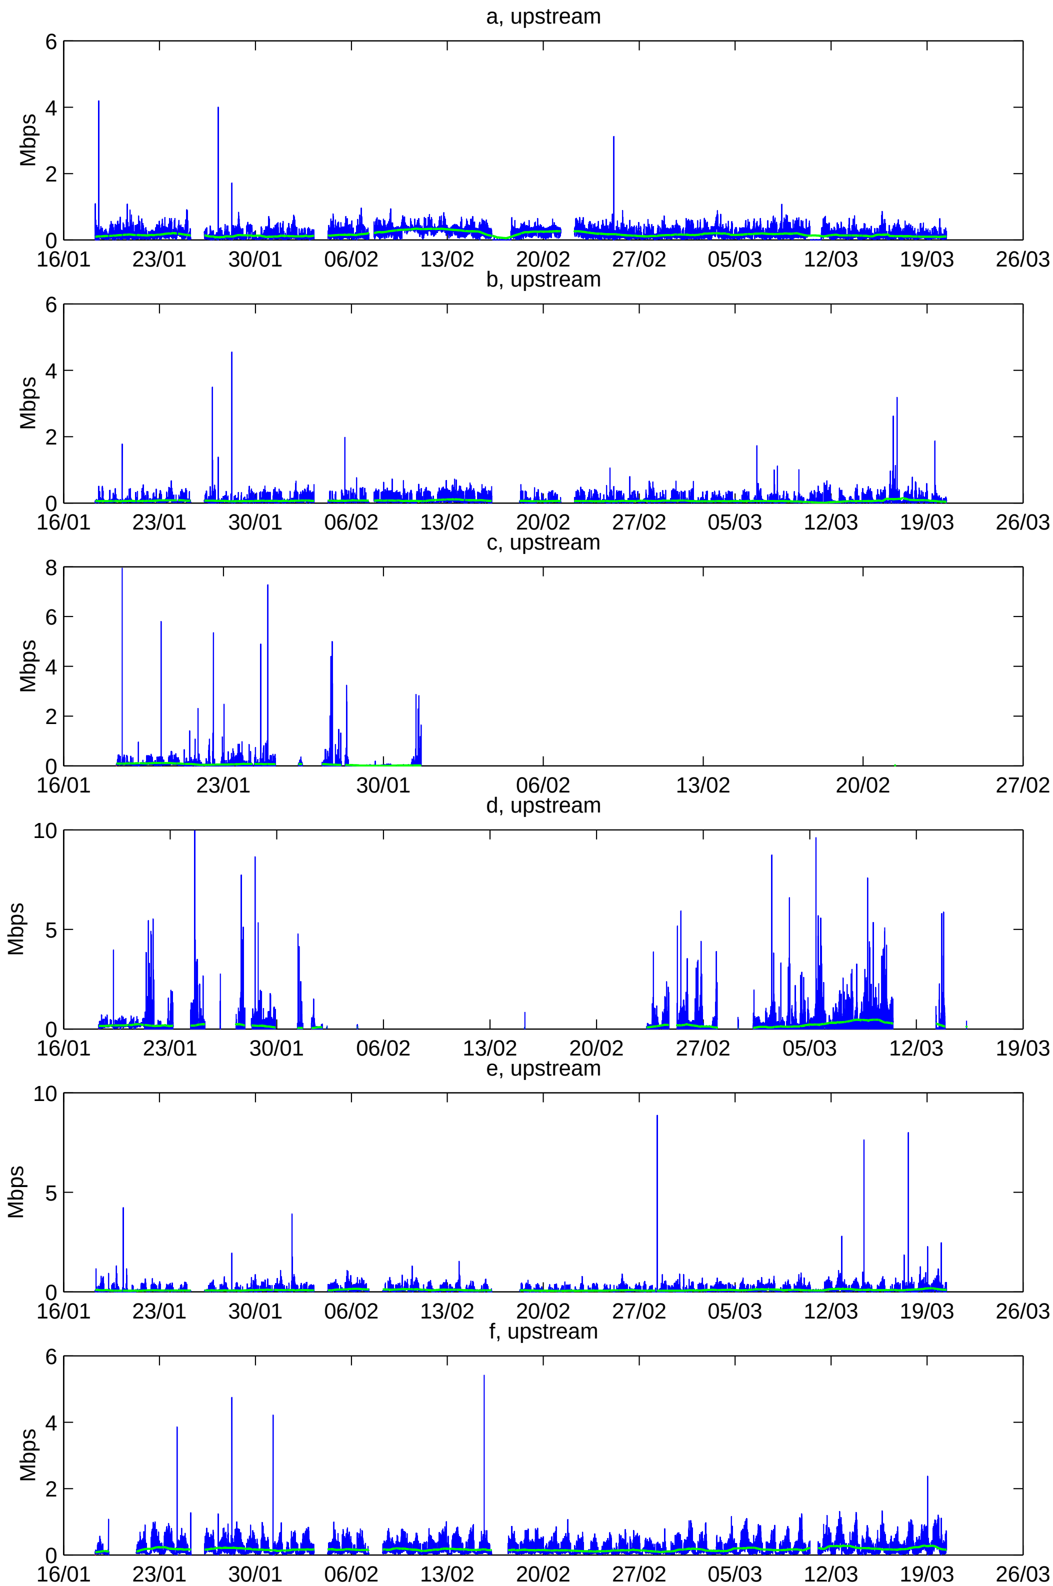

Fig. 3. Upstream CMTS traffic. The solid line shows one minute average rates (in Mbps) and the pale line shows a two week moving average of the data rate.

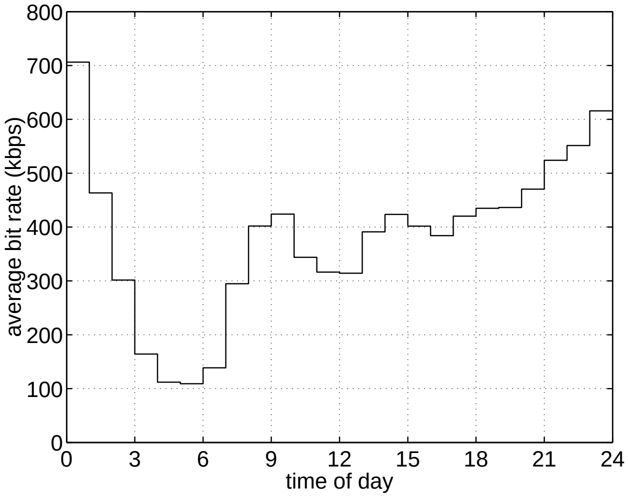

(a) dataset a

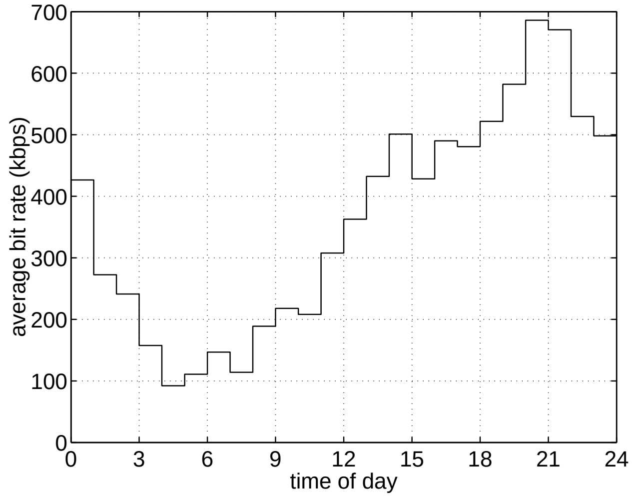

(b) dataset b

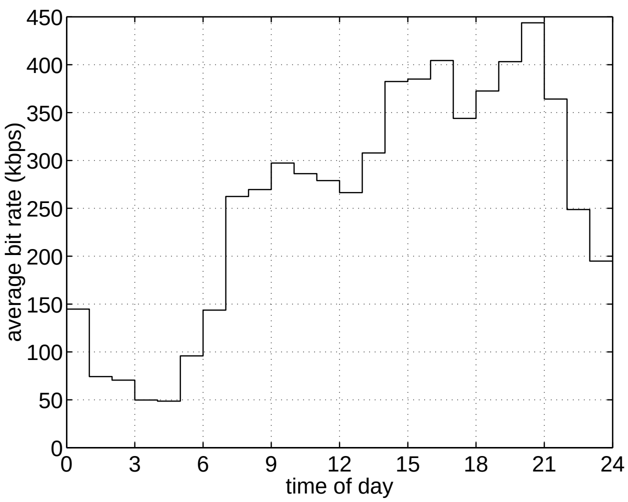

(c) dataset e

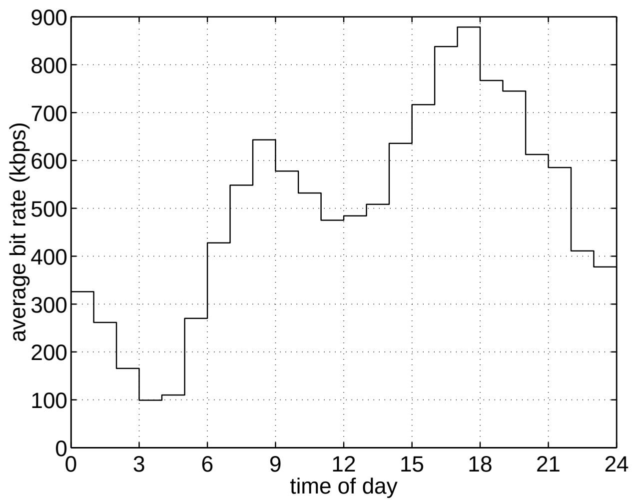

(d) dataset f

Fig. 4. Average traffic by time of day on each of the four investigated CMTSs.

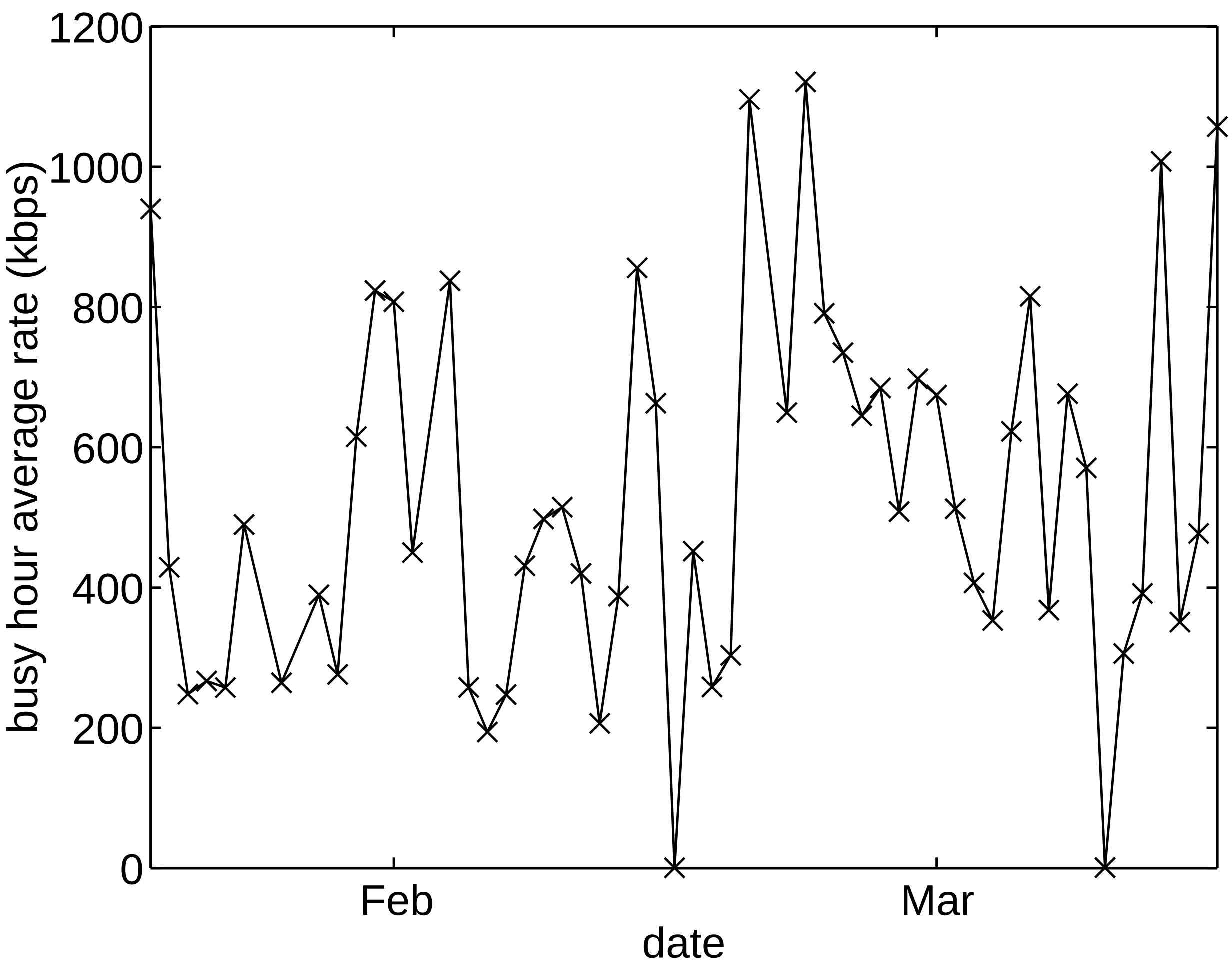

Fig. 5. 1 hour average of the traffic on a during the busy hour of each day.

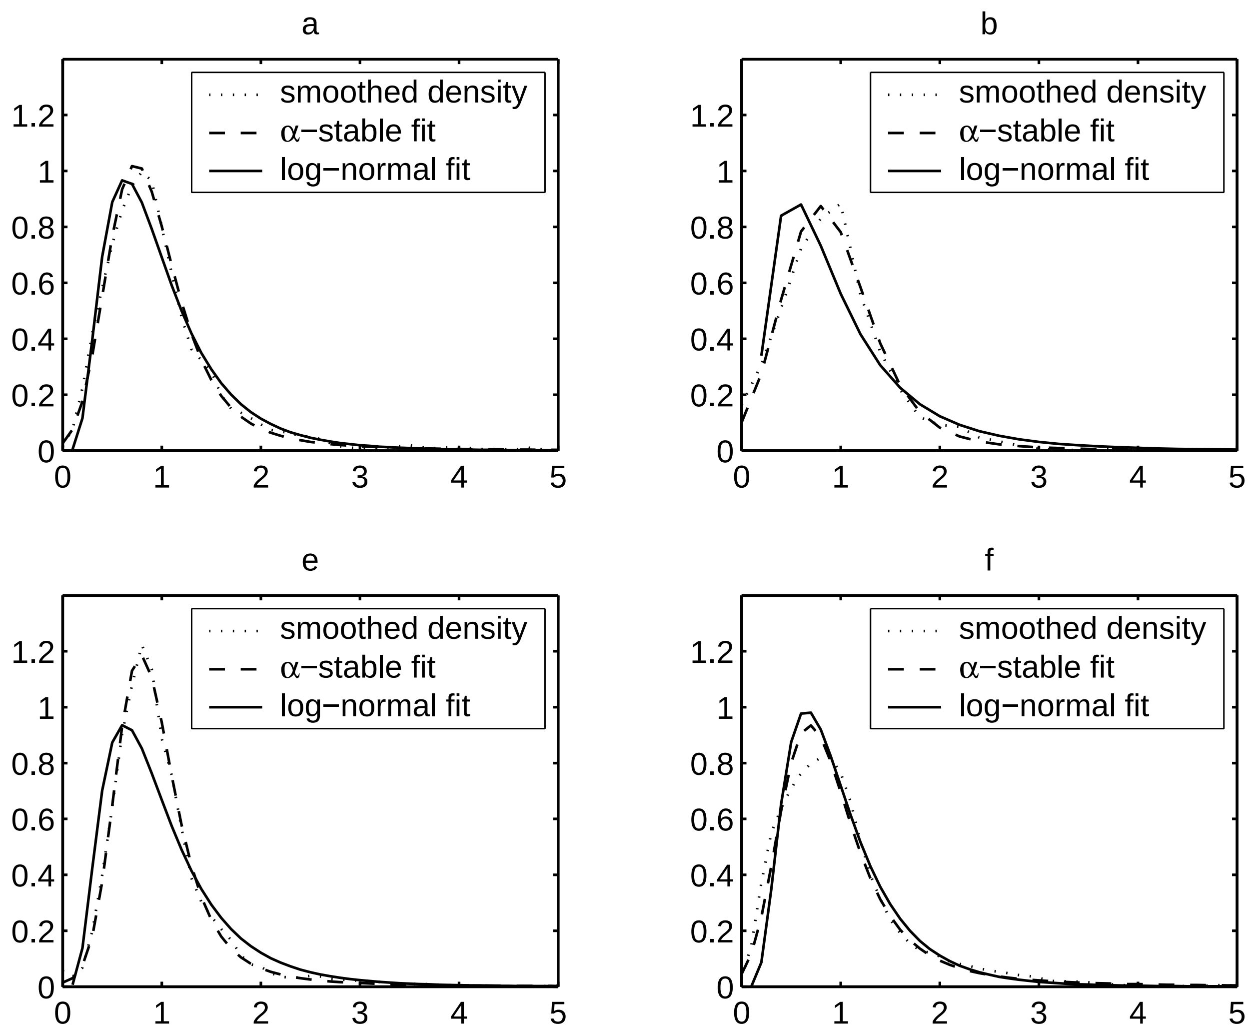

Fig. 6. The smoothed probability density of the normalized marginal rates (dotted), a log-normal fit to the data (solid) and an α-stable fit to the data (dashed).

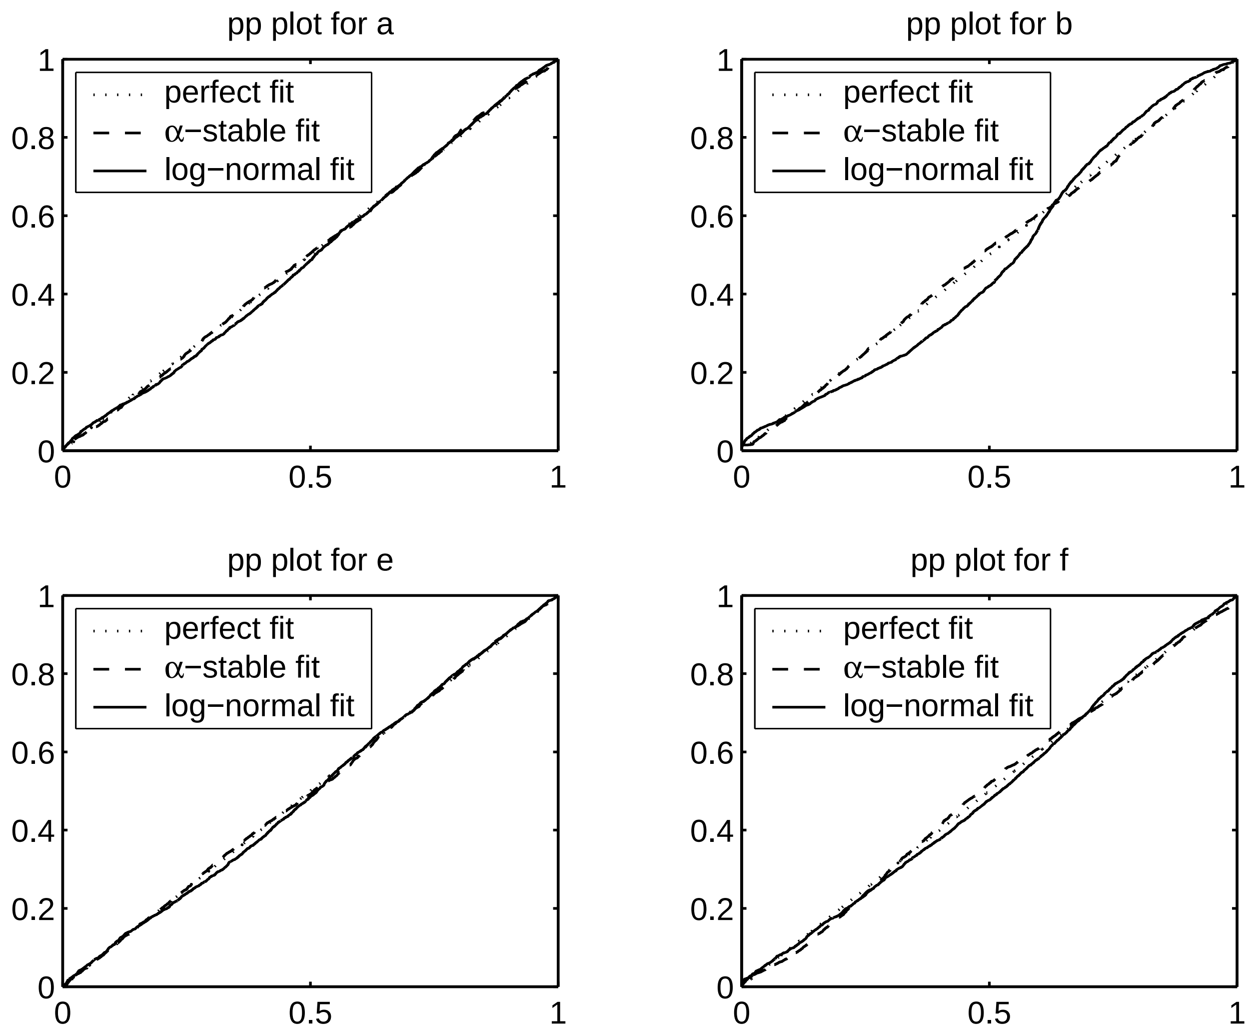

Fig. 7. pp-plots showing the quality of the log-normal and α-stable fits.

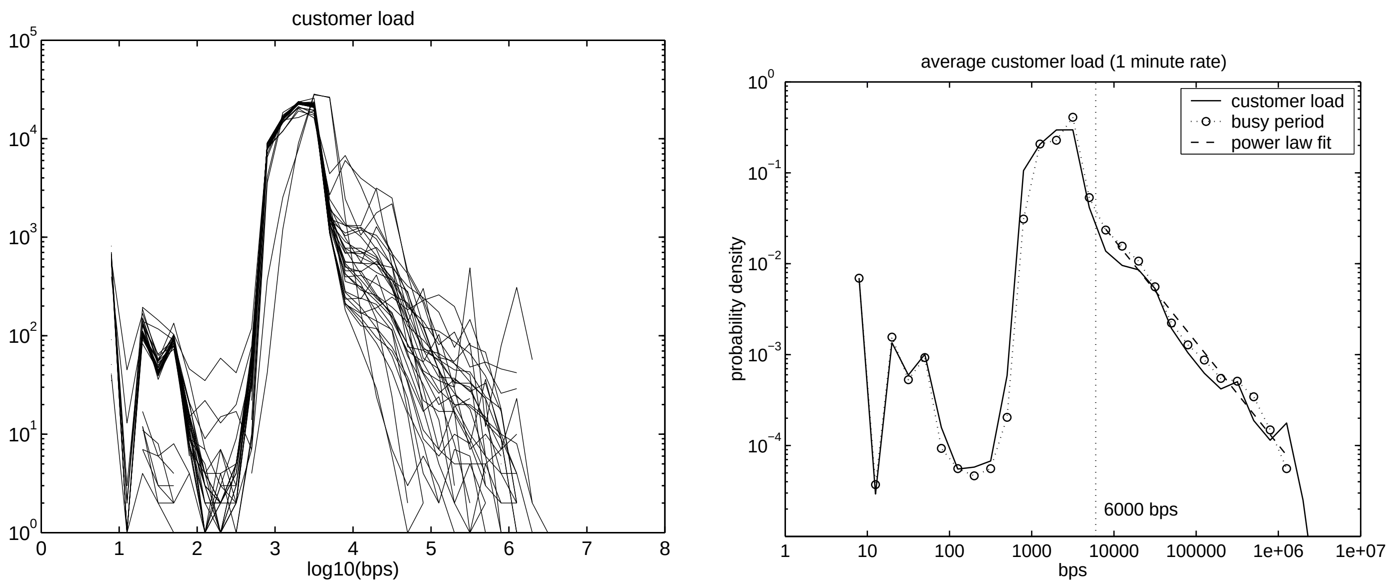

(a) individual marginals (b) average marginals

Fig. 8. Plots of the marginal distributions of the individual cable modems sampled from data set (a). Note the log-log axis.

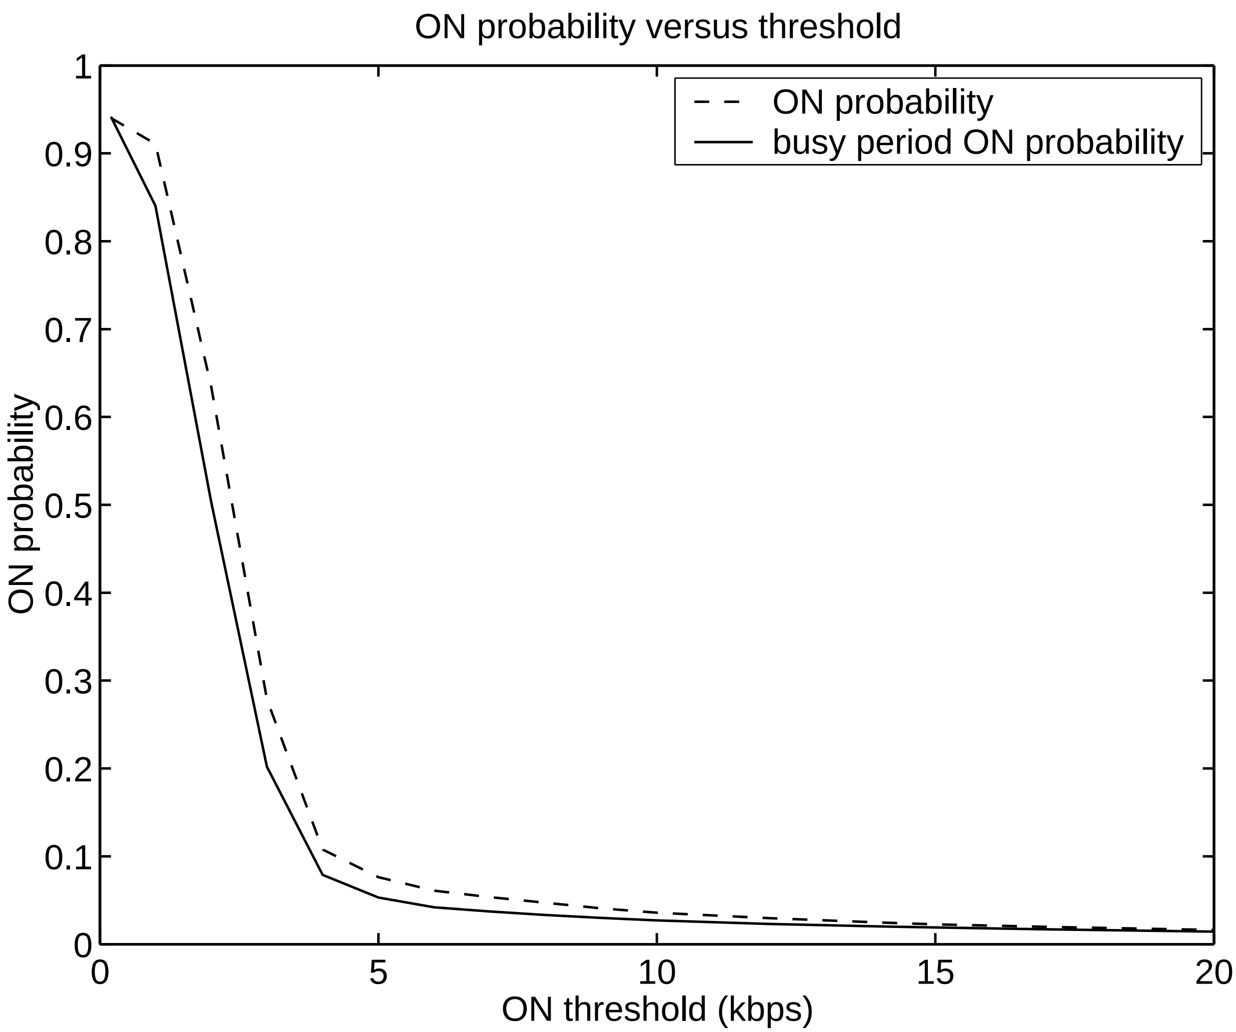

Fig. 9. The effect of changing the threshold on the ON probability for data set (a).

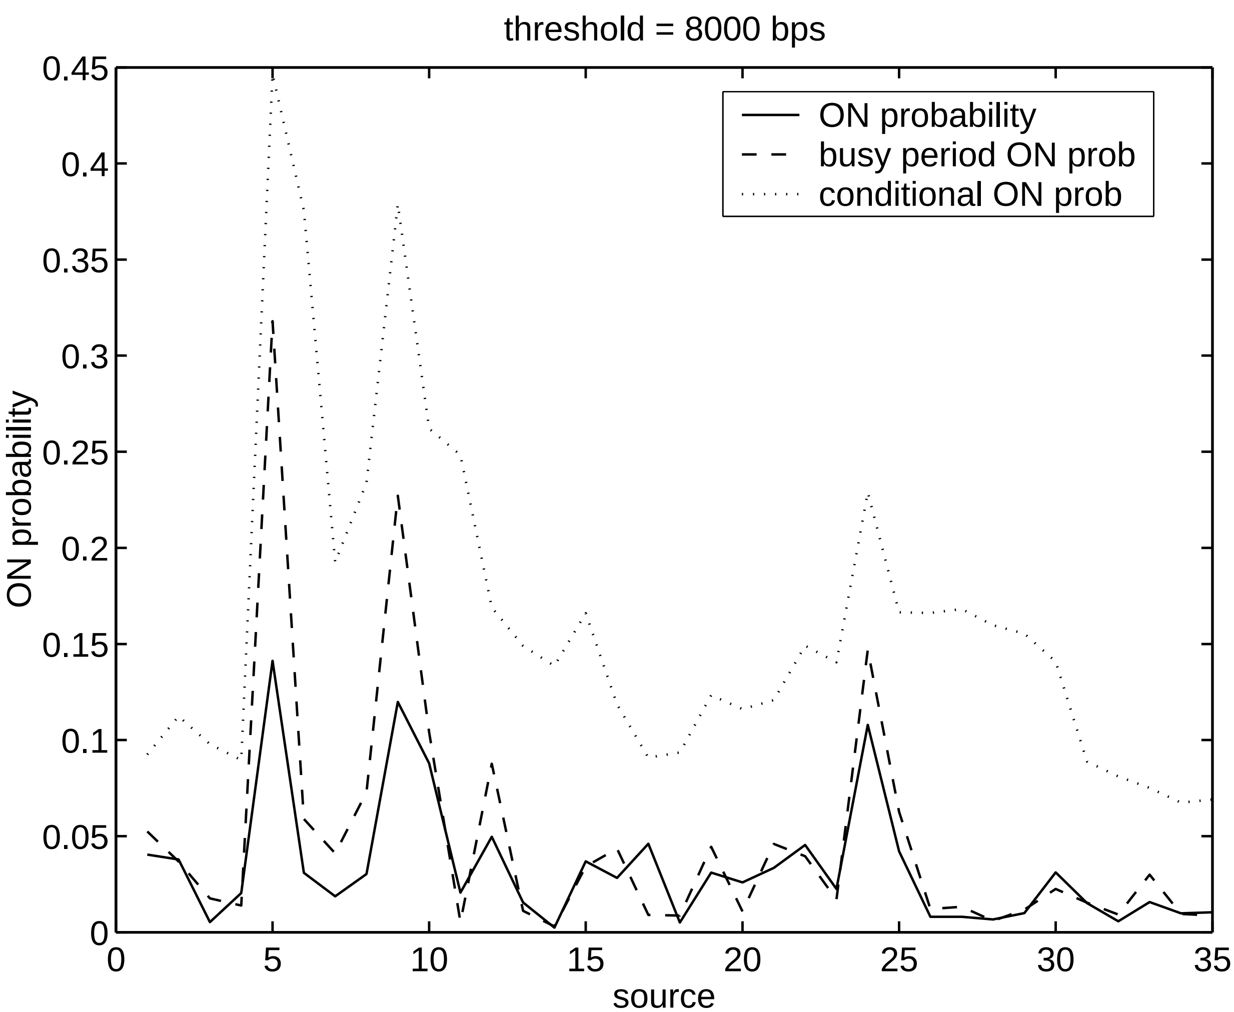

Fig. 10. The ON probability per source for data set (a): the linear interpolation between sources is for visualization purposes only, as the source number is arbitrary.

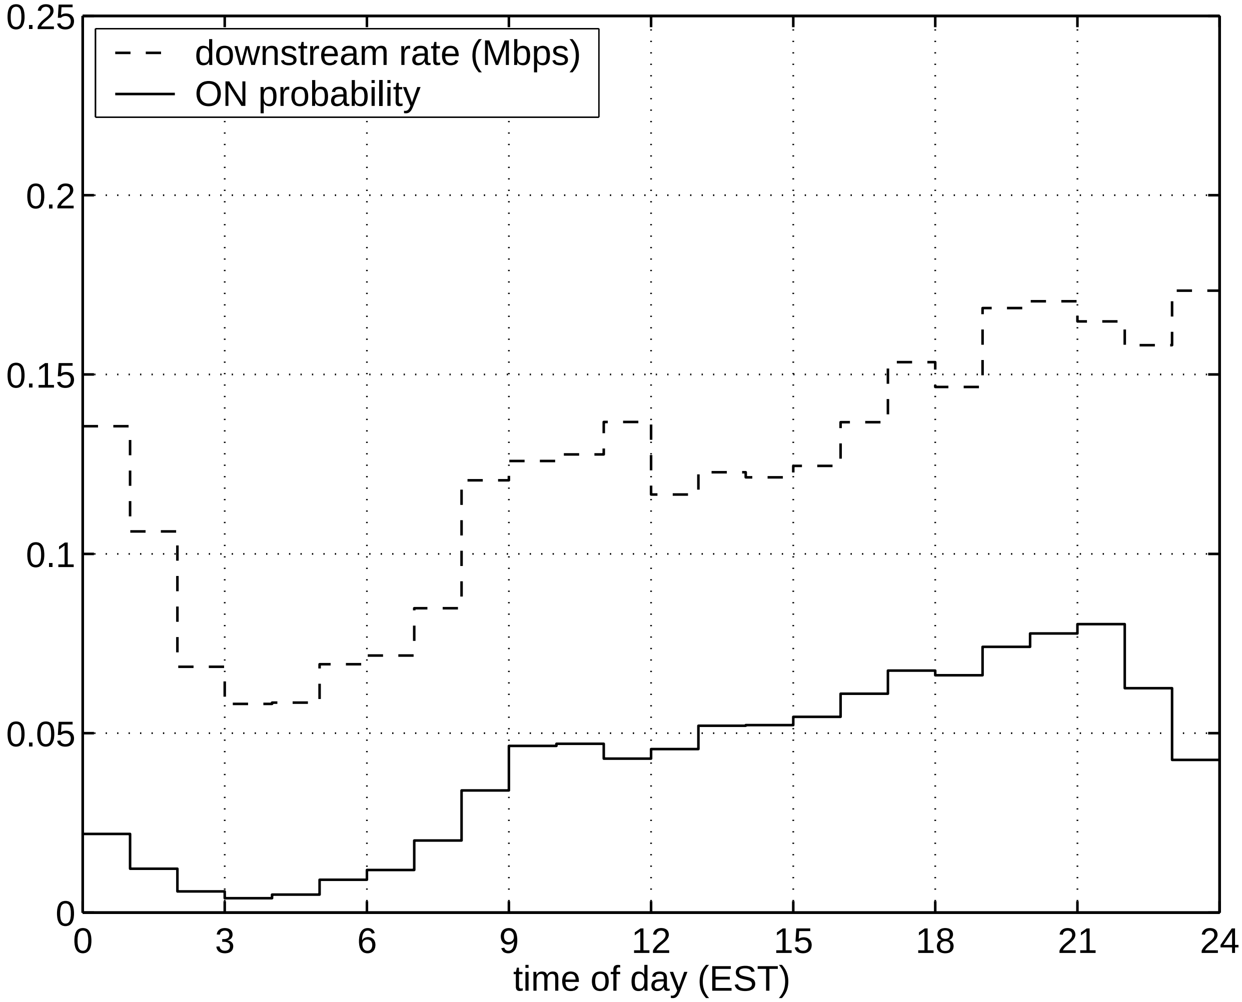

Fig. 11. The ON probability, and average rate for the sampled modems by time of day for data set (a).

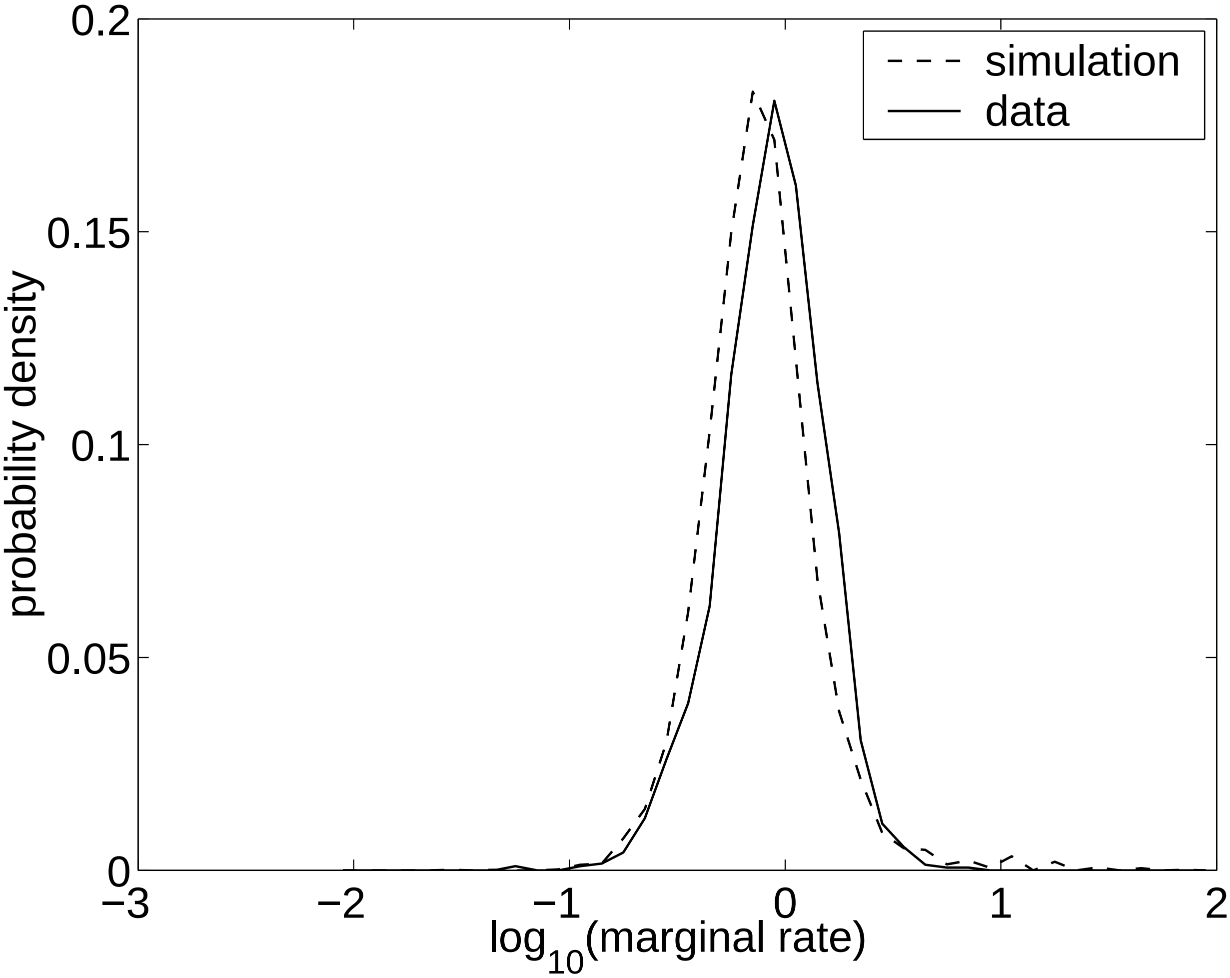

Fig. 12. A comparison of the simulated, and measured normalized marginal probability densities for data set (a).

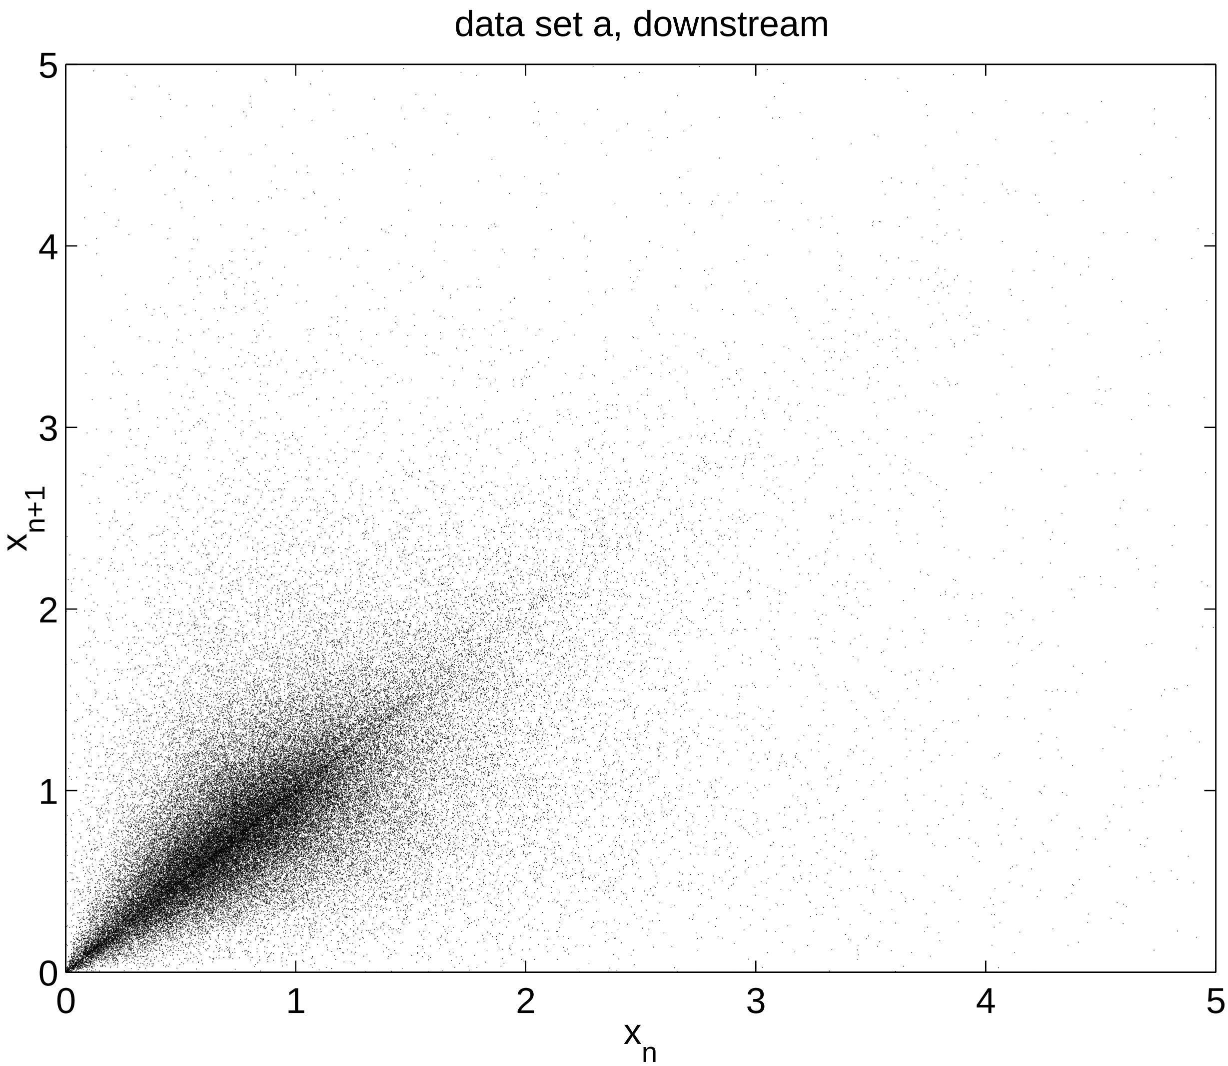

Fig. 13. A plot showing correlations in the renormalized data.