

ESTIMATING EXTREME VALUES IN DELAY-MEASUREMENTS VIA PROBES

ERIC VAN DEN BERG AND JOERG ROTHENBUEHLER

ABSTRACT. In the present study we analyze end-to-end probe packet delays statistically, motivated by the emergence of Service Level Agreements on network performance parameters. We use parsimonious parametric models from Extreme Value Theory, enabling us to make predictions on future delays at the upper end of and outside the range of the available data. In particular, we are able to model the maximum delay seen over a given time period (e.g. 1 hour) accurately with a Generalized Extreme Value distribution. This raises the possibility of including guarantees on maximum delays of the kind: "a delay of x ms will not be exceeded in more than one hour per month", which are akin to the requirements on which Extreme Value Engineering is based.

1. INTRODUCTION

Increasingly, Internet Service Providers are offering Quality of Service (QoS) guarantees for services exceeding the common 'best-effort' type [18, 20, 7], such as Virtual Private Networks (VPNs) or extranets. Upon subscribing to a VPN or extranet service, a customer is provided with a contract outlining the service quality to be delivered, the Service Level Agreement (SLA).

In order to determine whether SLA requirements are being met, careful monitoring is necessary. In fact, besides a set of metrics, an SLA should contain for each metric: a precise definition, the acceptable value range, and the measurement technique. Since SLA compliance relies in part on measurements, both from the customer's and from the service provider's point of view, it is important to have an idea of the statistical fluctuations to expect in these measurements. This enables the service provider to include accurate, attainable QoS guarantees in the SLA. A common performance metric included in SLAs is (one-way) packet

Date: January 17, 2003. *Keywords:* Traffic and performance measurements, Traffic Management .

©Telcordia Technologies, Inc. 2002. Address Evdb (Corresponding author): Telcordia Technologies, Inc., 445 South Street, Morristown, NJ 07960. Email: evdb@research.telcordia.com. Address JR: School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY 14853. Email: jrothenb@orie.cornell.edu.

delay [18, 7, 20]. This metric is often measured via periodic injection of probe packets. The metric for delay specified in the SLA is a function of the measured delay of these packets, for instance the sample average [20]. If the sample average is taken over a large enough number of packets, or timescale, then close predictions and promises can be made about its value by virtue of the central limit theorem, see e.g [20]. On smaller time-scales however, the sample average probe delay often fluctuates relatively widely about the average packet delay, e.g. because of the burstiness in traffic arrival processes. Hence, average performance guarantees are difficult to make on a short time scale, and are not a good reflection of customer experience.

A similar situation was encountered in local telephone central offices in the 1970s [3, 9]. Traffic engineering in the Bell System had traditionally been accomplished using measurements of average traffic loads during a time consistent busy hour. However, it was found that in small telephone offices only around 20% of the daily peak loads occurred in the time consistent busy hour. As an alternative, the Extreme Value Engineering (EVE) methodology used the distribution of the hourly peak loads with new service criteria, that better reflected the experience of a customer in the system. Engineering was based on a 20 (business) day return period load. A return period load is the load which will be exceeded only once, on the average, in an interval of time, the return period. Faced with highly varying loads due to internet connections and an increase in telecommuting, some local telephone companies have recently taken a second look at EVE to engineer the switching equipment in their central offices. If Extreme Value Engineering is an effective way to engineer the network to reflect good customer experience, the corresponding performance criteria, such as a return period packet delay, may yield effective performance criteria to include in an SLA.

In this paper we consider one-way delays measured for inserted probe packets as a potential SLA metric. We conduct a statistical analysis of end-to-end probe packet delays collected in Advanced Networks & Services' Surveyor project. Using parsimonious models from Extreme Value theory we are able to make predictions on future high delays, possibly outside the range of the available data. In particular, we accurately predict the maximum delay seen over a given time period (e.g. 1 hour) with a Generalized Extreme

Value distribution. Thus, it appears feasible to include in an SLA performance guarantees on maximum delays of the kind: "a delay of x ms will not be exceeded in more than one hour per month".

Extreme value analysis has long been used as a tool in insurance, and is gaining popularity in risk management in the financial industry [8, 14, 12]. To our knowledge, this is the first application of extreme value analysis to internet performance data. However, in the context of SLAs, we can view this as a similar application: how to set an appropriate delay guarantee, so that the risk of SLA violations is minimized.

A very recent paper discussing delay distributions on fixed internet paths is [10], based on one-way-delay measurements taken in the RIPE project in Europe, analyzed in [4]. The RIPE project collects data very similar to the data collected in the Surveyor project, which we analyze here. In fact, for a comparison of these two and other end-to-end performance measurement projects, see e.g. [6].

The remainder of the paper is organized as follows: In Section 2, we describe the nature of the datasets used in this study. In Section 3, we give a brief introduction to Extreme Value theory. In Section 4, we describe the results of our data analysis, and finally we present our conclusions in Section 5.

2. THE SURVEYOR PROBE DATA

The Surveyor project [15], led by Advanced Networks & Services, Inc., provides a measurement infrastructure that currently measures end-to-end unidirectional delay, loss, and routing among a set of about 50 measurement probes at higher education and research sites throughout the Internet [11]. Surveyor measures the One-way Delay [1] and One-way Loss [2] metrics developed by the Internet Protocol Performance Metrics (IPPM) working group of the IETF. This has the benefit of measuring (potential) standard metrics. Such standard metrics are also good candidates to be used in SLA specification.

Past studies [16] have shown that many Internet paths are asymmetric. In these circumstances, round-trip latency measurements e.g. via the 'ping' command, measure the performance of two different paths.

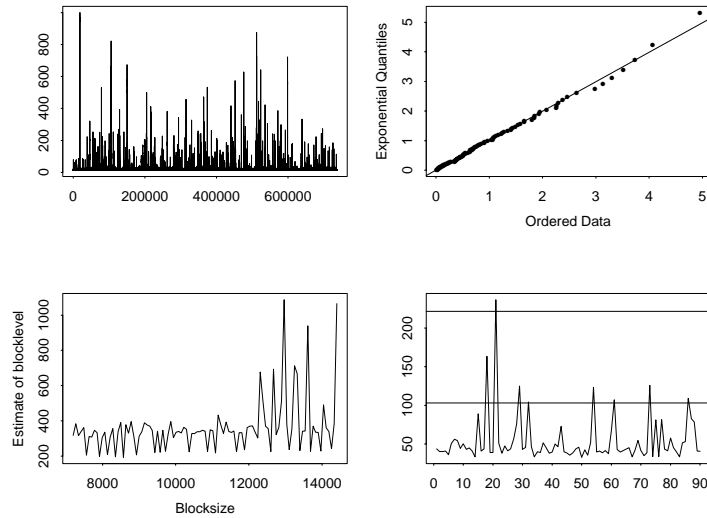


FIGURE 1. Probe packet delays (in ms). Top: University of Wisconsin to Harvard University (1999). Bottom: University of Colorado to Harvard University (2001). See text for explanations.

However, even if the path is symmetric, network load may be quite different in the two directions. One-way measurements allow us to measure these different entities separately, creating a clearer performance picture.

An example dataset of One-way Delay measurements between GPS monitors at the University of Wisconsin and Harvard University, taken from May 21, 1999 to May 28, 1999, is shown in Figure 1, upper left. More details on the measurement methodology will be given below. A fuller account can be found in [11].

In Surveyor, delay and loss are both measured using the same stream of injected test packets. These packets are scheduled according to a Poisson process simulated by a pseudo-random number generator running on the sending machine, at an average rate of 2 packets per second. This rate is chosen to keep impact of the test traffic on aggregate traffic minimal, while collecting as many measurements as possible. An additional frequency-limiting factor is the disk space needed to store the measurements. Each test packet is only 40 bytes. It consists of a sequence number and a timestamp, totaling 12 bytes, which are transported using UDP. The Surveyor measurement machines are equipped with GPS hardware for time synchronization. Therefore, timestamps carried by test packets have global meaning, and packet delay is computed simply

by subtracting the time in the received packet from the current time. Since the receiver has a copy of the same random number generator, initialized with the same seed, it can estimate the sender's timestamps very accurately. If a packet does not arrive within 10 seconds, it is considered to be lost. Lost packets are treated as sent packets with infinite delay by Surveyor. In this paper, however, we consider delays of non-lost packets.

IP packets experience delay from three distinct sources: transmission delay, propagation delay and queuing delay. On the links measured in this project (almost all T3 or faster) transmission delay is almost negligible compared to propagation delay. Furthermore, as long as IP and underlying layer 2 routing remains constant for a given path, both transmission and propagation delay are constant. The minimum observed source-destination delay in a given timeperiod, e.g. a minute, is close to the propagation delay [11], hence step functions in the delay often indicate routing changes. Daily traffic reports for the Surveyor project, including some summary statistics can be found at [15].

We had access to measured probe delays (measured in milliseconds) for 72 source-destination monitor pairs, measured from 20:00:00 Friday, May 21, 1999 to 20:00:00 Friday, May 28, 1999. Many time series of the probe delays exhibited step functions in the minimum delays, indicative of routing changes as mentioned above. For our statistical analysis, we purposely avoided these datasets. In this paper, we present the analysis of two datasets from the 72 source-destination pairs, which were representative of a large fraction of the total, while not suffering from obvious routing 'flaps' or excessive nonstationarity. In addition, we obtained one dataset of measurements taken two years later, from May 21, 2001 to May 28, 2001. The datasets are for the source-destination pairs University of Wisconsin - Harvard (1999), University of Colorado - Harvard (1999) and University of Colorado - Harvard (2001).

For the statistical analysis, we only considered weekday data. This eliminated the obviously different weekend measurements for University of Wisconsin - Harvard, as well as the small effect of the one route change on Sunday for that dataset.

3. EXTREME VALUE ANALYSIS: THEORY AND STATISTICAL TECHNIQUES

3.1. Extreme Value Theory. In this section we give a brief introduction to Extreme Value Theory. An excellent reference on this subject is [8]. Extreme Value Theory describes the fluctuations of

$$M_n := \max\{X_1, \dots, X_n\}, \quad (3.1)$$

where $\{X_1, \dots, X_n\}$ are i.i.d. random variables with cumulative distribution function F . By our assumptions,

$$P(M_n \leq x) = P(X_1 \leq x, \dots, X_n \leq x) = F^n(x).$$

Let x_F denote the (possibly infinite) right endpoint of F . That is, $x_F = \sup\{x \in \mathbb{R} : F(x) < 1\}$. Then one can show that M_n converges to x_F , with probability 1, as n approaches ∞ . This can be compared to the situation described by the strong law of large numbers. Let $\bar{X}_n := \frac{1}{n}S_n := \frac{1}{n}\sum_{i=1}^n X_i$. The sample mean \bar{X}_n of n i.i.d. random variables converges to the true population mean μ with probability one, as n approaches ∞ . The Central Limit Theorem further specifies the asymptotic distribution of \bar{X}_n . If (X_1, \dots, X_n) are i.i.d. random variables following a distribution with finite mean μ and variance $\sigma^2 < \infty$ then,

$$\frac{S_n - n\mu}{\sqrt{n}\sigma} \xrightarrow{d} Z \sim N(0, 1). \quad (3.2)$$

In a similar fashion, the following theorem describes the possible limit laws for $c_n^{-1}(M_n - d_n)$ for appropriate sequences c_n and d_n . It is the basis of Extreme Value Theory and its applications in this paper.

Theorem 3.1. [17, 8] *Let (X_n) be a sequence of i.i.d. random variables with common distribution F . Let $M_n := \max\{X_1, \dots, X_n\}$. If there exist constants $c_n > 0$ and $d_n \in \mathbb{R}$, and a non-degenerate distribution function H , such that*

$$c_n^{-1}(M_n - d_n) \xrightarrow{d} H \quad (3.3)$$

then H belongs to the following parametric family:

$$H_\xi = \begin{cases} \exp(-(1 + \xi x)^{-1/\xi}) & \xi \neq 0, \\ \exp(-\exp(-x)) & \xi = 0, \end{cases} \quad (3.4)$$

where $1 + \xi x > 0$. By introducing a location parameter μ and a scaling parameter ψ we obtain a model family flexible enough to allow fitting datasets of maxima, as explained later. This three parameter model family is called the Generalized Extreme Value Distribution, GEV:

Definition 1. (*Generalized Extreme Value Distribution*)

$$H_{\xi, \mu, \psi} = \begin{cases} \exp(-(1 + \xi(\frac{x-\mu}{\psi}))^{-1/\xi}) & \xi \neq 0, \\ \exp(-\exp(-(\frac{x-\mu}{\psi}))) & \xi = 0, \end{cases} \quad (3.5)$$

where $1 + \xi(\frac{x-\mu}{\psi}) > 0$. In relation to Theorem 3.1 and the above definition, the following questions arise:

- (i) Given an extreme value distribution H_ξ , what conditions on the distribution function F guarantee that the normalized maxima M_n converge weakly to H_ξ ?
- (ii) Can different constants $\{c_n > 0\}$ and $\{d_n\}$ imply convergence to different limit distributions H ?

Fortunately, the answer to the last question is No, thanks to the Convergence of Types Theorem, see e.g. [17].

Therefore, the following definition makes sense: If (3.3) holds for X_i with distribution F , and some extreme value distribution H_ξ , we say that F is in the *Maximum Domain of Attraction of H_ξ* , and write $F \in MDA(H_\xi)$.

To answer (i), we have to distinguish between three different cases: $\xi > 0$, $\xi < 0$ and $\xi = 0$. The MDA's for these three cases have been worked out, along with the normalization constants for the most common distributions. In the following we give a very brief overview. A more detailed and formal discussion, including on how to choose the normalizing constants $c_n > 0$, and $d_n \in \mathbb{R}$, such that $c_n^{-1}(M_n - d_n) \xrightarrow{d} H$, can be found in [8] in sections 3.3 and 3.4.

3.1.1. *Domain of attraction for $\xi > 0$.* A distribution function F belongs to the Maximum Domain of Attraction of $H_{\xi, \xi > 0}$, if and only if

$$\bar{F}(x) := 1 - F(x) \sim x^{-\alpha} L(x) \quad (3.6)$$

as $x \rightarrow \infty$ for some slowly varying function L , where $\alpha = 1/\xi$. A function L is called *slowly varying* if $\lim_{t \rightarrow \infty} \frac{L(tx)}{L(t)} = 1$. Functions $\bar{F} = 1 - F$ that satisfy (3.6) are called *regularly varying with index $-\alpha$* , and the class of regularly varying functions with index $-\alpha$ is denoted by $RV_{-\alpha}$. The Cauchy, Pareto, Burr and the Stable Distribution with exponent $\alpha < 2$ are examples of distributions in $MDA(H_{\xi})$.

3.1.2. *Domain of attraction for $\xi < 0$.* The most important fact about the maximum domain of attraction of $H_{\xi, \xi < 0}$ is, that all its members have a finite right endpoint $x_F < \infty$. Well known distributions in $MDA(H_{\xi, \xi < 0})$ include the Uniform distribution and the Beta distribution.

3.1.3. *Domain of attractions for $\xi = 0$.* $MDA(H_0)$ contains most distributions with an infinite right endpoint that have light right tails. We say that a distribution has a light right tail, if all moments $\mathbb{E}[(X^+)^k]$ exist and are finite. This is in contrast to the distributions in $MDA(H_{\xi, \xi > 0})$, which have only finite moments up to order $\alpha = 1/\xi$. H_0 itself has the property that $1 - H_0(x) \simeq e^{-x}$. This indicates that all distributions with an exponential or a ‘‘close to’’ exponential tail are in $MDA(H_0)$. In particular the Exponential, Gamma, Normal and Log-normal distribution all belong to $MDA(H_0)$.

3.2. Statistical methods for Extremes.

3.2.1. *Maximum Likelihood Estimation for the GEV.* If we wish to fit the GEV to a dataset, we have to consider serial dependence in the data. In [13] it is shown, that under a certain technical condition, referred to as $D(u_n)$, this dependence only affects the scaling parameters μ and ψ , but not ξ . $D(u_n)$ states a specific kind of asymptotic independence. It can be interpreted as stating that $\{X_i\}$ should not have a strong long range dependence. For a more formal discussion see [13] or [8], section 4.4.

We proceed as follows: We divide the data set in blocks of same sample size. In each block we determine

the maximum. The set of thus obtained block wise maxima is treated as an i.i.d. sample from GEV. This last assumption of independence of the block maxima is justified if the block size is chosen large enough, by the particular asymptotic independence implied by Condition $D(u_n)$.

This procedure raises the question: How many observations should make one block? On the one hand we have to make sure that blocks are large enough, so that their maxima are really i.i.d. and their distribution is close enough to a GEV. On the other hand, we want to keep the block size as small as possible, to obtain a sufficiently large number of maxima. We usually tried several different block sizes and then checked the quality of the fit to determine an applicable block size.

The three parameters of the GEV are estimated using Maximum Likelihood Estimation. A numerical procedure is needed to find the solutions to the complex likelihood equations. We used EVIS Version 3 with Splus 5.0 to carry out the calculations. If $\xi > -0.5$, [19] proves, that the MLEs are consistent and asymptotically efficient estimators. That is, they are asymptotically Normal distributed and their Covariance matrix is the inverse of the Fisher-Information matrix.

The goodness of the fit can be tested using the following transformation:

$$X \sim H_{\xi, \mu, \psi} \implies Y = \Phi^{-1}(H(X)) \sim N(0, 1) \quad (3.7)$$

where Φ denotes the standard normal cdf. To test whether the model actually fits the data, we transform the block-wise maxima using the estimated parameters. Those transformed maxima are then tested for normality with a Kolmogorov-Smirnov or a Chi-Square test of fit. If we fail to reject the null hypothesis that the transformed data is normal, we conclude that the estimates provide a satisfactory fit of the GEV to the block wise maxima. Additionally we may look at QQ-plots and similar exploratory tools.

3.2.2. k-Block Return Period Level Estimation. In the following section we show how Extreme Value Theory can be used for a candidate SLA guarantee. Our main tool is the *k-block* (return period) *level*.

Definition 2. *The k -block level is the value $b_{n,k}$ with the following property:*

$$\mathbb{P}[M_n > b_{n,k}] = 1/k, \quad (3.8)$$

where $M_n = \max(X_1, \dots, X_n)$. If M_n represent hourly maxima then $b_{n,k}$ is the level we expect to be exceeded only in one out of k hours on average. If M_n follows a GEV, we see that:

$$b_{n,k} = H_{\xi, \mu, \psi}^{-1}(1 - 1/k) = \mu - \frac{\psi}{\xi} \left(1 - (-\log(1 - 1/k))^{-\xi}\right). \quad (3.9)$$

We can therefore obtain estimates of blocklevels that may be outside of the data range, using a fitted GEV model by replacing ξ, μ, ψ with the corresponding MLE $(\hat{\xi}, \hat{\mu}, \hat{\psi})$. It is also possible to obtain asymptotical confidence intervals as outlined in [14].

4. EXTREME VALUE ANALYSIS OF PROBE DATA

The datasets considered in this section are probe packet data collected as described in Chapter 2.

4.1. Estimating Extreme Delays and Return Levels. We worked with block sizes of 7200, 10800 and 14400 for each dataset. These sizes correspond to time intervals of 1 hour, 1.5 hours and 2 hours, respectively. Other block sizes could of course also be considered. Block sizes smaller than 7200 usually resulted in bad fits and dependent block wise maxima while choosing larger block sizes resulted in too few data points to obtain a reliable fit.

4.1.1. *University of Wisconsin - Harvard.*

The block-wise maxima used in each of the three estimations were tested for independence. The test involved the methods of Turning-Points and Difference-Sign (see [5], p 312-313), as well as an examination of the ACF plots. We found no indication of dependence among the maxima. Furthermore, both Goodness-of-fit tests and QQ-plots (see Fig. 1, upper right) failed to detect any significant inadequacy of the fitted

TABLE 1. Estimates of the three parameters of the GEV for three different choices of blocksize.

Blocksize	7200		10800		14400	
Number of Blocks	103		69		52	
Estimate of ξ (95% C.I.)	.30	(.12 .49)	.34	(.08 .59)	.21	(-.06 .49)
Estimate of μ (95% C.I.)	149	(127 171)	194	(163 224)	230	(186 274)
Estimate of ψ (95% C.I.)	98	(80 117)	108	(82 135)	136	(101 172)

model. The parameter estimates in Table 1 indicate that the distribution of the delays is heavy tailed. The tail index $\alpha = 1/\xi$ seems to be between 3 and 5. The confidence interval for the block-size of 14400 is very wide due to the small sample size. Both other confidence intervals do not contain zero, which lends support to a hypothesis of heavy tailed delays. We present our estimates of the blocklevels, based on blocksize of 7200, in Table 2. The table lists blocklevels, the estimates and the corresponding C.I. in the first three columns. The remaining three columns address the accuracy of these estimations. They show the number of block maxima that we expect to exceed the corresponding block level in a data set of the same sample size, together with a confidence interval. These numbers, listed in 'Expected Exceedances' and '95% C.I.', are compared with the observed number of block maxima that exceed the estimated block level, listed under 'Observed Exceedances'. If the estimates are accurate, the number of observed exceedances should roughly equal the number of expected exceedances and fall in the confidence interval around our estimate.

We see from Table 2, that the number of observed and expected exceedances are roughly equal for each level. We find no evidence against the assumption that our estimates are accurate.

4.1.2. *University of Colorado - Harvard 1999.* We obtained the parameter estimates given in Table 3.

TABLE 2. Estimated k-block levels, confidence intervals. See the text for detailed explanations.

Level	Estimate	95% C.I.	Expected Exceedances	95% C.I.	Observed Exceedances
$b_{103,20}$	622	(498 875)	5.15	(2 8)	7
$b_{103,50}$	882	(662 1418)	2.06	(0 4)	1
$b_{103,100}$	1132	(792 2031)	1.03	(0 2)	0
$b_{103,200}$	1439	(925 2899)	.51	(0 1)	0

TABLE 3. Estimates of the three parameters of the GEV for three different choices of blocksize.

Blocksize	7200		10800		14400	
Number of Blocks	108		72		54	
Estimate of ξ (95% C.I.)	-.34	(-.48 -.19)	-.4	(-.60 -.19)	-0.43	(-.68 -.19)
Estimate of μ (95% C.I.)	59	(51 69)	55	(44 66)	52	(40 64)
Estimate of ψ (95% C.I.)	197	(185 210)	221	(207 236)	235	(220 251)

Surprisingly, this second dataset does not seem to be heavy tailed. The estimated value of the shape parameter ξ is negative. All used diagnostics described above failed to detect any significant inadequacy of the fitted model. Table 4 presents the estimated block levels along with the backtesting analysis described in the previous section.

TABLE 4. Estimated k-block levels, confidence intervals. See the text for detailed explanations.

Level	Estimate	95% C.I.	Expected Exceedances	95% C.I.	Observed Exceedances
$b_{108,20}$	309	(299 332)	5.40	(3 8)	8
$b_{108,50}$	327	(315 356)	2.16	(0 4)	3
$b_{108,100}$	337	(322 371)	1.08	(0 2)	0
$b_{108,200}$	345	(327 386)	0.54	(0 2)	0

Since the data is not heavy tailed, the confidence intervals for the k -block levels are much narrower, compared to the other two data sets. Again, we did not find an indication of a bias in the k -block level estimates.

4.1.3. *Colorado - Harvard 2001*. We obtained the following parameter estimates:

TABLE 5. Estimates of the three parameters of the GEV for Harvard - Colorado 2001.

Blocksize	7200		10800		14400	
Number of Blocks	90		60		45	
Estimate of ξ (95% C.I.)	.58	(.37 .79)	.60	(0.36 0.83)	1.05	(.63 1.48)
Estimate of μ (95% C.I.)	7.9	(6 9.8)	10.2	(7.2 13.1)	8.4	(4.5 12.4)
Estimate of ψ (95% C.I.)	41	(39 43)	44	(41 47)	45	(42 48)

This third data set consists of delays of packets between the same two monitors as the previous dataset, collected two years later. This time the delays appear to be heavy tailed, as the estimates of the shape parameter are all positive and the corresponding confidence intervals exclude 0. The estimates of the shape parameter based on block sizes of 7200 and 10800 are very close, while the one based on the blocksize of 14400 has a completely different value. A more detailed analysis confirmed that $\hat{\xi}=.6$ is the estimate resulting from most choices of a blocksize. The variability of the estimates increases for large block sizes, due to small sample sizes. Estimates based on those large block sizes are not reliable. While both the Turning-Points and Difference-Sign test don't detect any dependence, an examination of the ACF plots show some dependence among the maxima coming from a blocksize of 7200. While both the Kolmogorov-Smirnov and the Chi-square Goodness of Fit Test fail to reject the Null-hypothesis of a good fit for the maxima coming from block sizes of 10800, the Kolmogorov-Smirnov Test rejects with a p-value of 0.03 for the blocksize of 7200. However, the estimates of the block levels based on the two different fits are very similar. Studying plots of estimates of block levels (see Fig. 1, bottom left, for 200-block level), we saw

that estimates based on smaller block sizes are very consistent, while using large block sizes may result in unreliable estimates. It became obvious that the main factor that influences the estimate of the block level is the estimate of ξ . We therefore proceeded using the fit based on a block size of 7200 to estimate the block levels despite the indicated problems. This way we may still interpret the maxima as hourly maxima. In Fig. 1, bottom right, we have plotted the hourly maxima with the 20- and 100-block levels.

TABLE 6. Estimated k-block levels, confidence intervals. See the text for detailed explanations.

Level	Estimate	95% C.I.	Expected Exceedances	95% C.I.	Observed Exceedances
$b_{90,20}$	103	(81 154)	4.50	(2 7)	8
$b_{90,50}$	157	(110 291)	1.80	(0 4)	2
$b_{90,100}$	221	(138 463)	0.90	(0 2)	1
$b_{90,200}$	317	(176 534)	0.45	(0 1)	0

5. CONCLUSIONS

In this paper we considered one-way delays measured for inserted probe packets as a potential SLA metric. Guided by the bursty nature of the observed probe packet delays, and the success of extreme value engineering for switches under highly fluctuating traffic, we apply extreme value analysis techniques to the estimation of large delays. We conducted a statistical analysis of end-to-end probe packet delays collected by Advanced Networks & Services' Surveyor project. Using parsimonious models from Extreme Value theory to estimate high delays, we were able to make predictions on future large delays, possibly outside the range of the available data. In particular, the maximum delay seen over a given time period (e.g. 1 hour) was accurately modeled with a Generalized Extreme Value distribution. Thus, it appears feasible to include in an SLA performance guarantees on maximum delays of the kind: "a delay of x ms will not be exceeded in more than one hour per month".

Acknowledgements We would like to thank Matt Zekauskas of Advanced Network & Services, Inc. for providing the network measurement data, Mark Garrett, K.R. Krishnan, Judith Jerkins and Arnold Neidhardt of Telcordia Technologies and Sidney Resnick and Gennady Samorodnitsky at Cornell University for useful discussions and comments on earlier versions of the manuscript.

REFERENCES

- [1] G. Almes, S. Kalidindi, and M. Zekauskas. A one-way delay metric for ippm. Request for Comments 2679, September 1999.
- [2] G. Almes, S. Kalidindi, and M. Zekauskas. A one-way loss metric for ippm. Request for Comments 2679, September 1999.
- [3] D.H. Barnes. Extreme value engineering of small switching offices. In *8th ITC*, Melbourne, Australia, 1976.
- [4] C.J. Bovy, H.T. Mertodimedjo, G. Hooghiemstra, H. Uijterwaal, and P. van Mieghem. Analysis of end-to-end delay measurements in internet. In *PAM*, Fort Collins, Colorado, March 2002.
- [5] P. Brockwell and R. Davis. *Time Series: Theory and Methods, 2nd edition*. Springer-Verlag, New York, 1991.
- [6] L. Cottrell. <http://www.slac.stanford.edu/comp/net/wan-mon/iepm-cf.html>, July 1999.
- [7] ANX eBusiness. <http://www.anx.com/downloads/SLA.pdf>.
- [8] P. Embrechts, C. Klüppelberg, and T. Mikosch. *Modelling Extremal Events for Insurance and Finance*. Springer, Berlin, 1997.
- [9] K.A. Friedman. Extreme value analysis techniques. In *9th International Teletraffic Congress*, Torremolinos, Spain, 1979.
- [10] G. Hooghiemstra and P. van Mieghem. Delay distributions on fixed internet paths. Preprint.
- [11] S. Kalidindi and M. Zekauskas. Surveyor: An infrastructure for internet performance measurements. In *INET*, San Jose, 1999.
- [12] S. Lauridsen. Estimation of value at risk by extreme value methods. *Extremes*, 3(2):107–144, 2000.
- [13] M.R. Leadbetter, G. Lindgren, and H. Rootzén. *Extremes and Related Properties of Random Sequences and Processes*. Springer, Berlin, 1983.
- [14] A. McNeil. Calculating quantile risk measures for financial return series using extreme value theory. Preprint., April 1998.
- [15] Advanced Networks and Inc. Services. Surveyor homepage. <http://www.advanced.org/surveyor>.
- [16] V. Paxson. *Measurement and Analysis of End-to-End Internet Dynamics*. PhD thesis, University of California, Berkeley, 1997.
- [17] S. Resnick. *Extreme Values, Regular Variation and Point Processes*. Springer-Verlag, New York, 1987.
- [18] SAVVIS. <http://www.savvis.net/customer/slas.html>.
- [19] R.L. Smith. Maximum likelihood estimation in a class of nonregular cases. *Biometrika*, 72:67–90, 1985.
- [20] UUNet. <http://www1.worldcom.com/uunet/terms/sla/>.