Delay Distributions on Fixed Internet Paths

Gerard Hooghiemstra and Piet Van Mieghem

Abstract— **Based on end-to-end delay measurements of IP probe-packets over a fixed path obtained from RIPE NCC, we model the stochastic part of the delay and propose statistical methods to analyse the data.**

The Internet traffic on the fixed path interferes with the IP probe-packets. This is modeled as an alternating on/off renewal process. On top of the delay caused by Internet traffic, the IP probe-packets experience a random processing delay due to scheduling and conversing of IP packets to various lower layer technologies in the routers on the fixed path.

The total delay is both modeled parametricly and by a non-parametric method. Although the data indicates that the end-to-end delay distribution is heavy tailed, neither a Pareto nor Weibull law provided sufficient accurate fits. The non-parametric method essentially enabled a stable deconvolution that led to the delay due to Internet traffic. Furthermore, the non-parametric method provides good estimates of the probability that the path is not loaded and bounds on the queueing tail probabilities. Especially these qualifiers are useful for deploying real-time services on Internet.

Keywords— **End-to-end delay, stochastic modeling.**

I. Introduction

Since the Internet is being regarded as the universal network for both non-real-time and real-time services, serious efforts are being devoted to verify whether the current best-effort Internet can satisfy certain end-to-end delay bounds for real-time service such as Voice over IP [18], [7], [21]. For example, a tolerable one-way mouth-to-ear delay for a voice communication [21] is about 150 ms, while the packetization delay depending on the codec varies from 20 ms to 80 ms, leaving a remaining network end-to-end delay budget ranging from 70 ms to 130 ms. In about 84% of the cases as shown below, that remaining network end-to-end delay is of the order of a one-way transmission time of an IP-packet between two (arbitrary) routers in Internet. Therefore, understanding the end-to-end delay components in the (Inter)network is crucial to assess the possibility and the level of quality of service (QoS) of real-time services and to improve the current Internet architecture.

Several papers [4], [16], [10] and [21] report end-to-end delay measurements, mostly based on round-trip time (RTT). Difficulties with clock synchronization, asymmetries in the one-way and return path and path-variations during the measurement limit RTT-based measurements. To surmount these problems and because large amounts of unicast traffic are necessary to cover substantial parts of the Internet much attention is payed to *multicast* inference (e.g. [2], [8] and [13]). The key idea in multicast inference is to obtain performance measures of common links in a multicast tree based on the statistics of the multicast users. However, the above measurement techniques do not provide the accuracy of active, one-way delay measurements as performed by RIPE NCC (outlined in Section II).

Delft University of Technology, Information Technology and Systems, P.O. Box 5031, 2600 GA Delft, The Netherlands. Email: G. Hooghiemstra@its.tudelft.nl and P.VanMieghem@its.tudelft.nl

The end-to-end delay D experienced by probe-packets over a fixed path containing h hops (routers) consists of two components: a deterministic D_d and a stochastic delay D_s. The deterministic delay D_d adds the contributions of the physical delay (speed of light/electromagnetic waves through the links and the routers, roughly $5\mu s/km$), the time between time-stamp generation and effective start of transmission of the probe-packet and of the bandwidth of a link (y b/s-links imply a delay of x/y s to transmit an x bit packet). The stochastic delay D_s is caused by interfering Internet traffic on that fixed path and by the random part of the processing delay generated by the operation of the h routers (table look-up, delay in the interface card, etc.). A detailed study of a single-hop delay is presented by Papagiannaki *et al.* [19]. In this article, we mainly focus on the stochastic delay D_s based on the RIPE NCC measurement data, while the deterministic delay D_d is discussed elsewhere [5]. Because of our confinement to separated, fixed paths, the stochastic delay D_s can be filtered satisfactorily from the total end-to-end delay D by subtracting the minimum end-to-end delay denoted by m experienced by the probe-packets during the day.

In Section II we first describe the RIPE NCC measurement configuration and present histograms of the total end-to-end delay D. In Section III, the stochastic delay D_T caused by Internet traffic (also called the queueing delay) is modeled by one or more renewal processes while the processing delay seems well modeled by a Gaussian. A statistical analysis on one particular, representative data set which outlines the method is given in Section IV. Results on various other paths by applying the analysis method of Section IV are presented in Section V. The appendix contains a mathematical proof.

II. Measurement data

RIPE NCC, the Network Coordination Centre of the Réseaux IP Européen, is continuously measuring the delay and the hopcount of IP-packets transmitted between fixed measurement boxes in some part of the Internet. At this moment (summer 2001), about 40 measurement boxes are scattered mainly over Europe. Between each pair of measurement boxes, small IP packets of a fixed length (100 bytes), called probe-packets, are transmitted with interarrival times of about 40 seconds, resulting in a total of about 2160 probe-packets per day. The sending measurement box generates an accurate time-stamp synchronized via GPS in each probe-packet, while the receiving measurement box reads the GPS-time of the probe-packet upon arrival. The end-to-end delay or one-way transit time of probe-packets is defined as the difference between these two time-stamps and has an accuracy of $10\mu s$. At regular times, path information (the number of hops, IP addresses of intermediate routers) between each pair of measurement boxes is obtained from the trace-route utility. The specific details of the RIPE NCC measurement configuration are described in [17].

In this article, we focus on end-to-end delay D of probe-packets on fixed paths between two measurement boxes. Specifically, by choosing a certain date of a day and two boxes, all end-to-end delays of probe-packets that follow precisely the same path between the boxes during that day are stored in a data set. From the large number of empirical delay distributions, the majority (over 80%) resembles the shape as illustrated in Figure 1, which we further coin as a *typical* delay distribution. Other, non-typical delay distributions are shown and discussed in a companion paper [5]. The total delay of probe-packets lies for the typical RIPE NCC measurements between 10 to 400 ms. Figure 3 shows that the average end-to-end delays lie around 20 ms to 40 ms. Non-typical delay distributions may broadly exceed 400ms with high probability making VoIP over these paths fairly impossible. In this article, only typical delay distributions will be analysed and modelled. To explain the various steps in the modeling, we will refer throughout this paper to the typical distribution of March 3, 2001 illustrated in Figure 1.

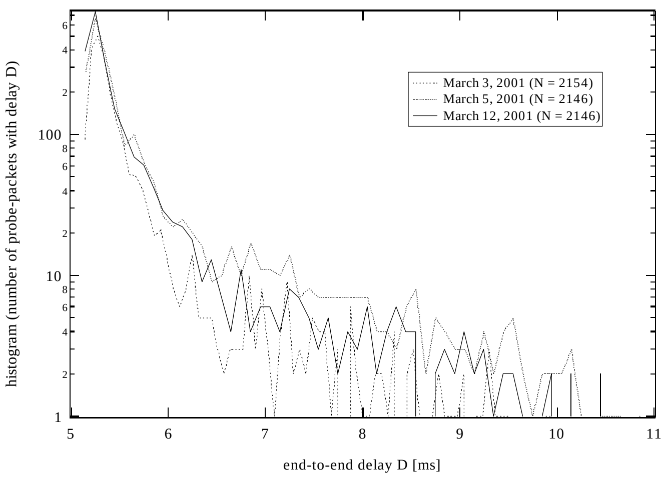

Fig. 1. The histogram of a typical delay, measured from box A to box B during a day. The number of samples N is found in the legend.

In this illustrative example, a total of 2154 probe-packets were transmitted over a fixed path between measurement box A and box B (the precise Internet Exchanges are confidential), showing a well-determined minimum transmission time of 5.11 ms, which justifies that the stochastic delay D_s is accurately obtained by subtracting this minimum transmission time from the end-to-end delay D. Furthermore, the histogram of the end-to-end delay D in Figure 1 has been plotted on a log-lin scale to show the remarkably long tail.

Since many articles report power law behavior for observables in Internet, we have analysed the data on log-log scale and found, indeed, that in an intermediate region, the normalized histogram (or the empirical probability density $f_D(x)$ of the end-to-end delay D) can be fitted reasonably well by a line of the form $\log f_D(x) = \beta_r \log x + \delta_r$, where the subscript refers to 'raw'. The fitted values of the power β_r (from an intermediate region only) are shown in Figure 2, together with the error bars.

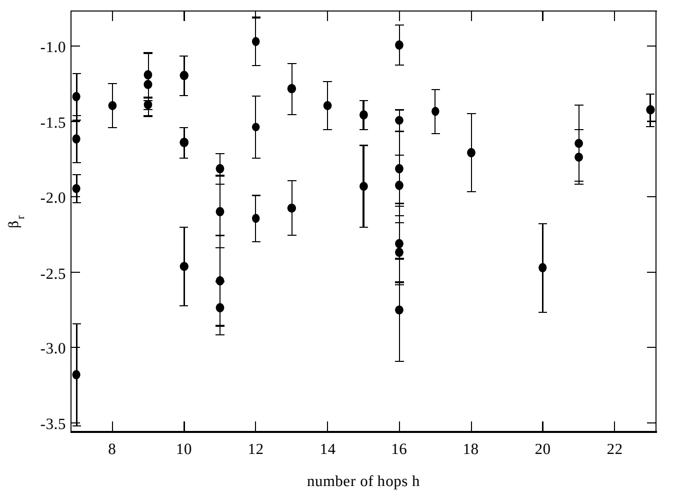

Fig. 2. The power β_r versus the hopcount of the fixed path for 38 typical paths, with average $\overline{\beta_r} = -1.78$ and standard deviation $s_{\beta_r} = 0.53$.

Figure 2 illustrates that the power β_r seems independent of the hopcount h (number of traversed routers in the fixed path). Furthermore, the logarithm of the same measured histogram $f_D(x)$ of the end-to-end delay D has been fitted by $\log(f_D(x)) = a_r |x - c_r|^{b_r} + k_r$ in some region where the fit eyed good. This type of fit inspires a Weibull-distribution rather than a polynomial distribution as suggested in [19]. We observed that c_r is close to the minimum delay D as illustrated in Figure 3. As also observed from Figure 1, the 'shape' parameter b_r was always negative (while $a_r > 0$) in discrepancy with the Weibull probability density function (11). Further, the parameter a_r and b_r seem correlated; a curve fit of b_r versus $\ln(a_r)$ suggests a linear correlation $b_r = -0.22 - 0.166 \ln a_r$. Again, it is very unlikely that the shape parameter b_r (nor a_r, c_r) is correlated with the hopcount h as illustrated in Figure 4.

The inspection of the 'raw' end-to-end delay data on various plotscales does not convincingly lead to insight in the tail behavior. An accurate estimate of the tail behavior or $\Pr[D > x]$ is desirable, especially in the deployment of real-time service in Internet. This motivates a more detailed study of the (stochastic) end-to-end delay $D(D_s)$.

III. Modeling the end-to-end stochastic delay over a fixed path.

At first glance, the stochastic end-to-end delay D_s over a fixed path of h hops may be modeled as a sum of the

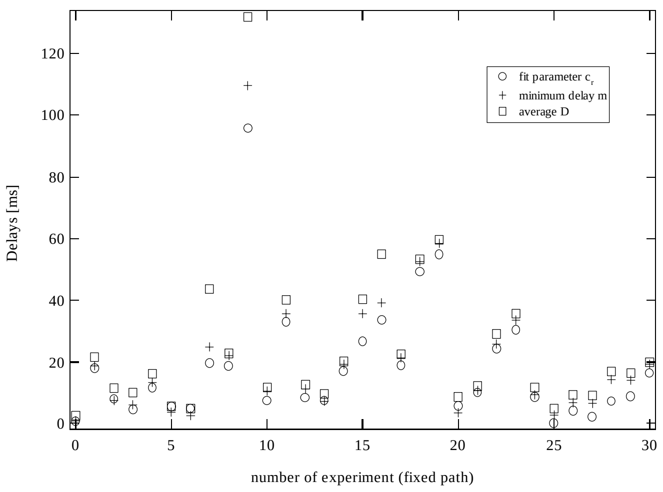

Fig. 3. The fit parameter c_r together with the minimum and average delay of the end-to-end delay D per measurement.

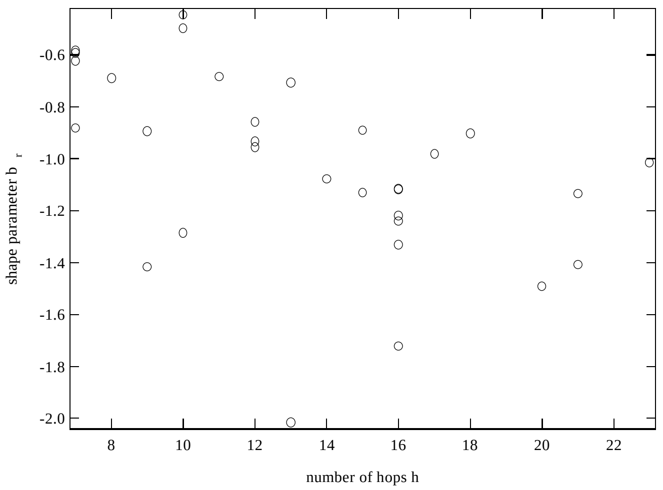

Fig. 4. The shape parameter b_r versus hopcount h for 31 paths with average $\overline{b_r} = -1.02$ and $s_{b_r} = 0.36$

stochastic delay $D_{s;j}$ incured per router j along that path,

$$D_s = \sum_{j=1}^{h} D_{s;j} \tag{1}$$

One may further assume to a good approximation that all $D_{s;j}$ are statistically independent such that the probability density function of D_s reduces to a convolution of the probability densities of these independent $D_{s;j}$. The delay per router $D_{s;j}$ can be described by various single server queueing models, with as simplest one an M/M/1-queue. This approach has been followed earlier in [20]. Although physically sound, when comparing the resulting probability density function derived from this model to the Internet measurements, two important discrepancies arose. First, the traffic intensity turned out to be very low and second, perhaps more interesting, the Internet data did not show a clear correlation with the hopcount h which contradicts the convolution model derived from (1). The seemingly independence of the end-to-end delay on the hopcount h as illustrated in Figure 2 and Figure 4 necessitates another approach. In fact, as discussed in Section III-C, we distinguish two components in the stochastic delay D_s: (a) the router processing delay which is additive as in (1) and (b) the delay caused by Internet traffic D_T which does not seem to reflect the additive structure.

In this article, we assume that the stochastic end-to-end delay of IP packets along a fixed path from a source to a destination is mainly caused by Internet traffic on or crossing the path. The idea is sketched in Figure 5 and is analogous to the delay experienced by travelling from A to B along a route with h intersections (or traffic lights). Indeed, depending on the cars (traffic) interfering with us and on the signs of the h traffic lights, the total time from A to B equals the end-to-end delay. The phenomenon of the disturbance by Internet traffic can be accurately modeled by renewal theory, which we present in this section. We model the lengths of epochs *without* Internet traffic on

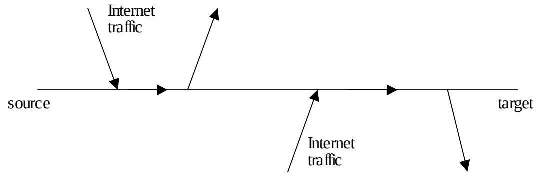

Fig. 5. The influence of Internet traffic

the fixed path by open times having a distribution function denoted by F and the lengths of epochs *with* Internet traffic by closure-times specified by distribution function G. During closure times the probe-packet experiences a delay untill the path opens, whereas during the open times the probe-packet travels without interference of other Internet traffic and incurs no stochastic delay. On top of the stochastic delay caused by Internet traffic, modeled by one or a concatenation of renewal processes, our probe-packet experiences additional stochastic delay caused by the scheduling of tasks within routers.

We will first compute the stochastic delay due to one Internet traffic stream, followed by a concatenation. Finally, these delays are augmented by the additional stochastic delay due to processing (scheduling and influence of lower layer technologies).

A. One renewal process

We consider two i.i.d. sequences: the sequence $X_1, X_2, \ldots$ representing the lengths of intervals during which our specific path is open (for the probe-packet) and the sequence $Y_1, Y_2, \ldots$ during which our specific path is closed (or blocked). The two sequences are also assumed to be independent of one another. The probability distribution function of X_j (Y_j) is denoted by F (G).

Now define on the postive real axis the renewal epochs: $X_1 + Y_1, \ldots, \sum_{i \le j}(X_i + Y_i), \ldots$, as indicated in Figure 6. A probe-packet that arrives at time t during a closed (Y)

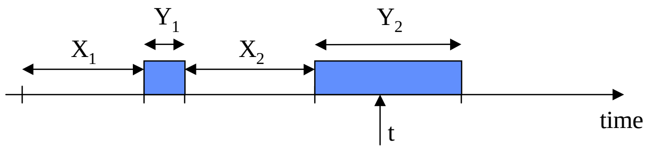

Fig. 6. An alternating renewal process: During the periods with length X_j the path is open; during the periods with lengths indicated by Y_j the path is closed.

interval experiences a *stochastic* delay of magnitude

$$\sum_{i\leq j}(X_i+Y_i)-t, \text{ if } t \in \left(\sum_{i\leq j-1}(X_i+Y_i)+X_j, \sum_{i\leq j}(X_i+Y_i)\right) \tag{2}$$

while probe-packets that arrive during open (X) intervals are not delayed by Internet traffic.

Let $N(t)$ be the number of renewals of $X+Y$,

$$N(t) = \max\{j \geq 0 : \sum_{i\leq j}(X_i+Y_i) \leq t\} \tag{3}$$

The stochastic delay β_t of a probe-packet arriving at time t due to Internet traffic then equals

$$\beta_t = 0,$$

if $\sum_{i=1}^{N(t)}(X_i+Y_i) < t < \sum_{i=1}^{N(t)}(X_i+Y_i) + X_{N(t)+1}$ and

$$\beta_t = \sum_{i=1}^{N(t)+1}(X_i+Y_i) - t,$$

if $\sum_{i=1}^{N(t)}(X_i+Y_i) + X_{N(t)+1} < t < \sum_{i=1}^{N(t)+1}(X_i+Y_i)$.

In the appendix we prove

Theorem 1: For non-arithmetic distributions F and G,

$$\lim_{t\to\infty} \Pr[\beta_t > u] = \frac{\int_u^\infty [1-G(x)]\,dx}{\mu_F + \mu_G}, \tag{4}$$

where $\mu_F = E[X] = \int_0^\infty [1-F(x)]\,dx$, is the average open time and $\mu_G = E[Y] = \int_0^\infty [1-G(x)]\,dx$ is the average closure time.

For the illustrative example in Figure 1, the total delay lies roughly between 5 and 11 ms. The stochastic delay covers roughly the interval from 0 to 6 ms. These values are small compared to the time between probe-packets, which is of the order of 40 s. Hence, if the Internet traffic on the path can be modeled by *an* alternating renewal process, it is reasonable to assume that we sample from the distribution given in Theorem 1.

The probability that the *stochastic* delay equals zero is

$$\begin{aligned} 1 - \lim_{t\to\infty} \Pr[\beta_t > 0] &= 1 - \frac{1}{\mu_F+\mu_G}\int_0^\infty [1-G(y)]\,dy \\ &= \frac{\mu_F}{\mu_F+\mu_G}, \end{aligned}$$

and is proportional to the mean durations μ_F and μ_G. If β denotes the steady state limit of the stochastic delay, then Theorem 1 can be rephrazed as

$$\begin{aligned} \Pr[\beta \leq u] &= 1 - \frac{\int_u^\infty [1-G(x)]\,dx}{\mu_F+\mu_G} \\ &= p + qG_R(u), \end{aligned} \tag{5}$$

where $G_R(u) = (\mu_G)^{-1}\int_0^u [1-G(x)]\,dx$ is the residual distribution of G and $p = \mu_F/(\mu_F+\mu_G)$, while $q = 1-p$. The form (5) exhibits the two components: the atom at $u=0$ with magnitude p, the probability that the probe-packet is not delayed, and a residual distribution function $qG_R(u)$, effective for $u > 0$.

B. More than one renewal process

In the preceding section the delay caused by Internet traffic was modeled by *one* alternating renewal process. It is more realistic to assume that different links of the path are used by various other Internet streams and that our probe-packet is possibly delayed by one or *more* Internet traffic streams. We assume that the different Internet-streams can be modeled as independent renewal processes as sketched in Figure 5.

Let us first concentrate on the stochastic delay modeled by two *independent* renewal processes, $\beta^{(1)} + \beta^{(2)}$. Assume that for $i = 1, 2$,

$$\Pr\left[\beta^{(i)} \leq u\right] = p_i + (1-p_i)G_R^{(i)}(u),$$

and where $\beta^{(1)}$ is taken statistically independent of $\beta^{(2)}$. We find that

$$\begin{aligned} \Pr\left[\beta^{(1)} + \beta^{(2)} \leq x\right] &= \int_0^x \Pr\left[\beta^{(2)} \leq x-u\right] d\Pr\left[\beta^{(1)} \leq u\right] \\ &= p_1 \Pr\left[\beta^{(2)} \leq x\right] + \\ &\quad \int_0^x (p_2 + q_2 G_R^{(2)}(x-u))\frac{q_1}{\mu_{G_1}}(1-G_1(u))\,du \\ &= p_1p_2 + p_1q_2G_R^{(2)}(x) + p_2q_1G_R^{(1)}(x) \\ &\quad + \frac{q_1q_2}{\mu_{G_1}\mu_{G_2}}\int_0^x G_R^{(2)}(x-u)(1-G_1(u))\,du. \end{aligned}$$

Hence the atom at 0 is p_1p_2 and the density part equals

$$\begin{aligned} &\frac{p_1q_2[1-G_2(x)]}{\mu_{G_2}} + \frac{p_2q_1[1-G_1(x)]}{\mu_{G_1}} \\ &+ \frac{q_1q_1}{\mu_{G_1}\mu_{G_2}}\int_0^x (1-G_2(x-u))(1-G_1(u))\,du. \end{aligned}$$

We observe that, for both one and two (indepedent) renewal processes, the stochastic delay leads to an atom at zero and a continuous part on $(0,\infty)$. By mathematical induction, the stochastic delay of a concatenation of three or more independent renewal processes remains a mixed distribution of the form (5).

C. Router processing delay

As mentioned in the introduction the delay D of a probe-packet consists of different components. In the previous subsection we have modeled the *stochastic* delay caused by Internet traffic. In this subsection we will derive a model for the probability density φ of the total end-to-end delay D, which also yields the density model for D_s by shifting the density φ over the minimum delay m towards zero.

As observed in [5], over 80 % of the delay histograms of the RIPE NCC data are Gamma-shaped (they are close to a Gamma probability density function, in first order). If the delay D consisted only of a deterministic component together with the random delay caused by Internet traffic, the RIPE data histograms would be identically equal to zero up to this deterministic length and then exhibit an atom at this position. This is clearly not the case and a reasonable explanation is as follows. Due to other computational tasks in the routers and the variable processing time which depends on specific layer technologies (IP over SDH/optics, IP over ATM, etc.), for each router a random (continuous) delay should be added on top of the stochastic delay caused by Internet traffic. In the terminology of queueing theory, the processing delay in each router can be interpreted as the (stochastic) service time of the router.

Figure 7 displays two histograms obtained from a simulation. In the left hand side histogram we have simulated the delay , which originates from one renewal process with exponential on and exponential off times shifted over $m = 5.11$. One clearly identifies the atom at 5.11 ms with magnitude $p = 9/16$. The right hand side histogram shows the convolution of the stochastic delay modeled by the same renewal process and a uniformly distributed processing delay, which indeed resembles a Gamma-shape.

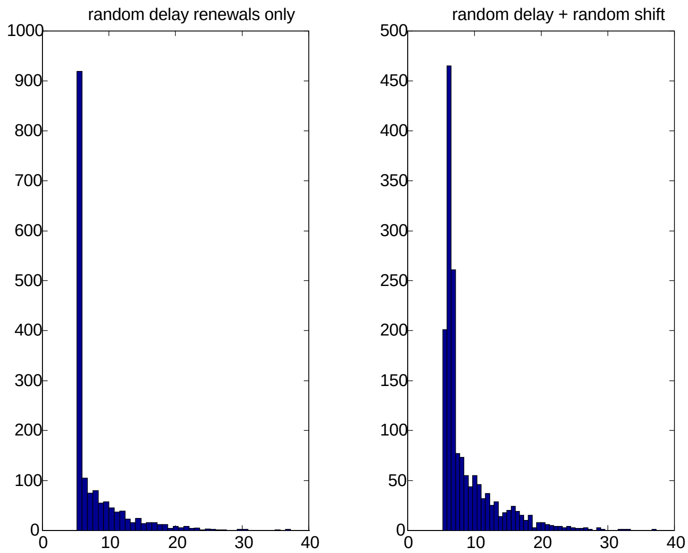

Fig. 7. Simulated delays (number of repetitions is 1598); left plot only renewal delays, right plot renewals plus a random shift

Summarizing, on top of the a stochastic delay caused by Internet traffic with distribution function of the form (5), the routers (all together) cause an additional 'processing' delay with probability density φ_2 which includes the minimum deterministic delay m. The independent sum of the minimum deterministic delay m, the Internet traffic delay and router processing delay has probability density function

$$\varphi(t) = p\varphi_2(t) + q\varphi_1 * \varphi_2(t), \quad t \geq 0, \tag{6}$$

with $\varphi_1 * \varphi_2(t) = \int_0^t \varphi_2(u)\varphi_1(t-u)\,du$.

The validity of the above approach can be further justified using the illustrative data set from the RIPE data and by lab-measurements on isolated routers. During this day a total of 2130 IP-packets[1] were sent over a fixed path from source box A to target box B. Analyzing the delay over periods of two hours the delay measurements during night-hours (from 2.00 AM to 8.00 AM) are found to be strikingly different from the remaining hours as illustrated in the table below

	0-2	2-4	4-6	6-8	8-10	10-12
mean delay (ms)	5.45	5.35	5.34	5.35	5.51	5.64
variance (ms)^2	0.21	0.05	0.03	0.04	0.41	0.57
minimum (ms)	5.17	5.14	5.15	5.14	5.11	5.17

	12-14	14-16	16-18	18-20	20-22	22-24
mean (ms)	5.76	5.83	5.59	5.65	5.57	5.49
var (ms)^2	0.80	0.85	0.50	0.44	0.37	0.26
min (ms)	5.18	5.18	5.15	5.18	5.18	5.19

The variance of the delay during the night ($< 0.05\ (\text{ms})^2$) is approximately 10 percent of the variance during the day. We therefore separated the delay measurements into a night and day period. The night period (from 2.00 AM-8.00 AM) contains 532 data-points, the remaining 1598, data-points (*not between 2.00 AM and 8.00 AM*) form the data for the day period. In Figure 8 we display the histograms of the total delay D.

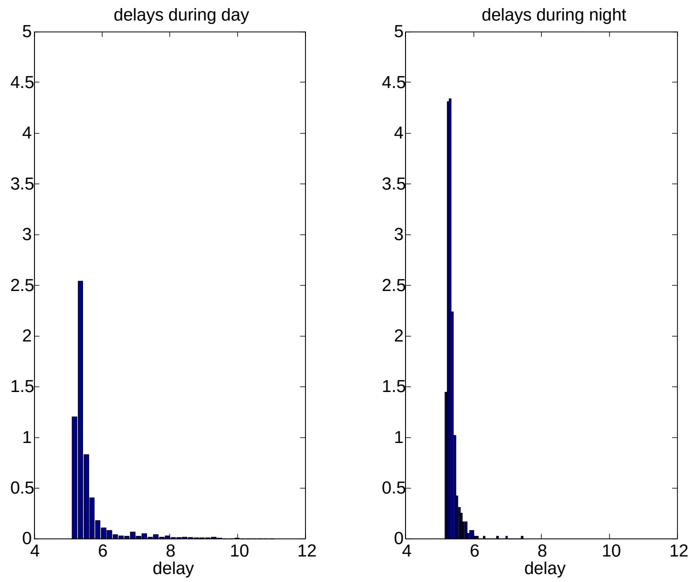

Fig. 8. Normalized histograms of delays; left during daytime right delay during night-time

During the night hours there is hardly any Internet traffic. Indeed, inspection of the (night-)histogram indicates

[1]We have removed 24 outliers whose value could be clearly shown as due to experimental errors.

that the density during the night hardly reflects Internet-delay, except perhaps for some values in the right tail beyond[2] 5.5 ms. If night-delays with values larger than 5.5 ms are ignored, a bell-shaped density remains that can be well approximated by a Gaussian density with the mean $\mu = 5.30$ ms and standard deviation $\sigma = 0.078$ ms.

The total delay measured with 1, 2 and 3 routers between two similar boxes as used in the RIPE data, but under lab conditions and on a physical negligible distance, gave similar results (cf. [5]). It is impossible to compare the mean router processing delay with the path mean $\mu = 5.30$, because the latter includes the unknown physical delay m. However the measured delay under lab conditions exhibits a symmetric density. Moreover, if we extrapolate the lab-measured standard deviation of $\sigma = 0.014$ ms for 1 router to the above path, containing $h = 10$ routers, we obtain $\sigma_h \approx \sigma\sqrt{h}$ or $\sigma_h \approx 0.044$ ms, which is of the same order of magnitude as our estimated value of 0.078 ms.

In summary the end-to-end delay D of a fixed path can be modeled by a density $\varphi(t)$ of the form (6). Apart from the delay caused by Internet traffic, which is modeled by the mixture $p+(1-p)\varphi_1(t)$, we identified that the processing delay (of all routers together) can be well approximated by a Gaussian density,

$$\varphi_2(t) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(t-\mu)^2}{2\sigma^2}} \tag{7}$$

where μ includes the minimum deterministic delay m.

IV. Statistical analysis of the end-to-end delay data.

In Willinger *et al* [22], an extensive study was made to asses the distributions of on- and off-periods for Internet traffic between fixed source and destination pairs. Quite convincingly they showed that the off-period follows a power law. However, from their log-log plot of the complementary distribution of the on-period, it is in our opinion less clear that this distribution exhibits a power law. Recall that the on-periods are the periods during which our probe-packet is blocked, if Internet traffic happens to use (partly) our fixed path.

The three proposed parametric models for the stochastic delay caused by Internet traffic D_T, (a) the exponential model, (b) the Weibull model and (c) the polynomial (or Pareto) model, all exhibit discrepancies with the data. Hence, parameter curve fitting or maximum likelihood estimation to unravel the Internet traffic from end-to-end delay measurements leads to poor tail results in the presence of a dominating processing delay component. Therefore, we have included a focused deep tail analysis while subsection IV-E concentrates on a deconvolution technique which yields a non-parametric maximum likelihood estimator (NPMLE) for D_T.

[2]The choice of the threshold 5.5 ms may seem arbitrary but can in fact be taken equal to the mean 5.34 ms of the night-data augmented with the difference between the mean 5.34 ms and the minimum 5.14 ms. The resulting threshold is then 5.54 ms.

A. An exponential density for $\varphi_1(t)$

In this subsection the delay D_T caused by Internet traffic is modeled by one alternating renewal process with exponential closure periods (the closure periods correspond to the periods where the Internet traffic has an on-period). Hence

$$\varphi_1(t) = \lambda e^{-\lambda t}, \qquad t \geq 0,$$

where λ^{-1} models the mean length of the closure period. The density (6) of the end-to-end delay can be evaluated exactly as

$$\varphi(t) = p\varphi_2(t) + \lambda(1-p)\int_{-\infty}^{t} e^{-\lambda(t-s)}\varphi_2(s)\,ds. \tag{8}$$

Assuming that μ and σ are known, the parameters λ and p can be estimated by the method of *maximum likelihood.* We define the log-likelihood L, given the observations $t_1, t_2, \ldots, t_n$,

$$\begin{aligned} L(\lambda, p) &= \log\left\{\prod_{i=1}^{n}\varphi(t_i)\right\} \\ &= \sum_{i=1}^{n}\log\left\{p\varphi_2(t_i) + \lambda(1-p)\int_{-\infty}^{t_i} e^{-\lambda(t_i-s)}\varphi_2(s)\,ds\right\} \end{aligned}$$

To find the arguments $\hat{\lambda}$ and $\hat{p}$ that maximize the likelihood or equivalently the log-likelihood we calculate the partial derivatives of L with respect to p and λ, and put them equal to 0. Using identity (8), these equations are equivalent to

$$\sum_{i=1}^{n}\frac{\varphi_2(t_i)}{\varphi(t_i)} = n, \tag{9}$$

$$\lambda^2\sum_{i=1}^{n}\frac{\int_{-\infty}^{t_i}(t_i-s)e^{-\lambda(t_i-s)}\varphi_2(s)\,ds}{\varphi(t_i)} = n. \tag{10}$$

Given the values of μ and σ in (7) the equations (9) and (10) can be solved numerically, to obtain $\lambda = \hat{\lambda}$ and $p = \hat{p}$, the ML-estimates.

We subsequently apply the model (8) to the realisations $t_1, t_2, \ldots, t_n$ of the $n = 1598$ delay-measurements between the boxes A and B, during the day period of March 3, 2001. The parameters μ and σ for the router processing delay in (7) are extracted from the night data, as indicated in Section III-C.

The ML-estimates obtained from (9) and (10) are

$$\hat{p} = 0.580 \quad \text{and} \quad \hat{\lambda} = 1.39$$

The fit of the density φ specified by (8) using the above estimates for μ, σ, p, λ, with the normalized histogram of the data is good, except in the tail as observed from the Figure 9.

B. A Weibull density for $\varphi_1(t)$.

A Weibull density is considered as model for $\varphi_1(t)$,

$$\begin{aligned} \varphi_1(t) &= abt^{b-1}e^{-at^b}, \qquad (11) \\ \int_0^t \varphi_1(u)du &= 1 - e^{-at^b}, \end{aligned}$$

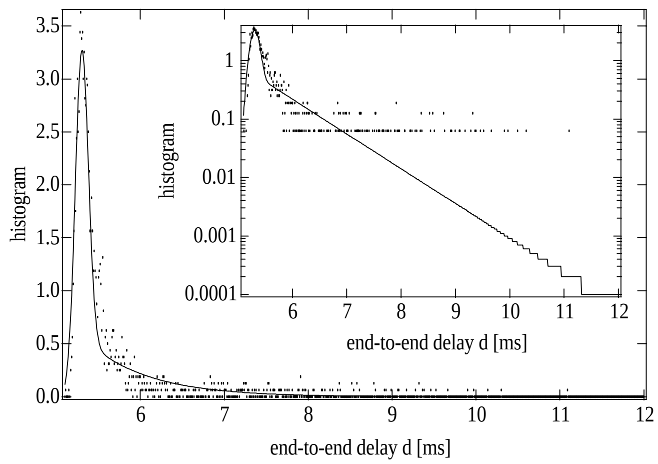

Fig. 9. Internet end-to-end measurements (dots) fitted with exponential closure times, i.e. model (8) in full line.

with expectation and variance

$$\begin{aligned}\int_0^\infty t\varphi_1(t)dt &= \frac{\Gamma\left(1+\frac{1}{b}\right)}{a^{1/b}} = \nu, \\ \int_0^\infty (t-\nu)^2\varphi_1(t)dt &= \frac{\Gamma\left(1+\frac{2}{b}\right)-\Gamma^2\left(1+\frac{1}{b}\right)}{a^{2/b}}.\end{aligned}$$

The ML optimization is best if $p = 0.30$, $a = 2.12$ and $b = 0.51$. Both fit and normalized histogram data are shown in Figure 10.

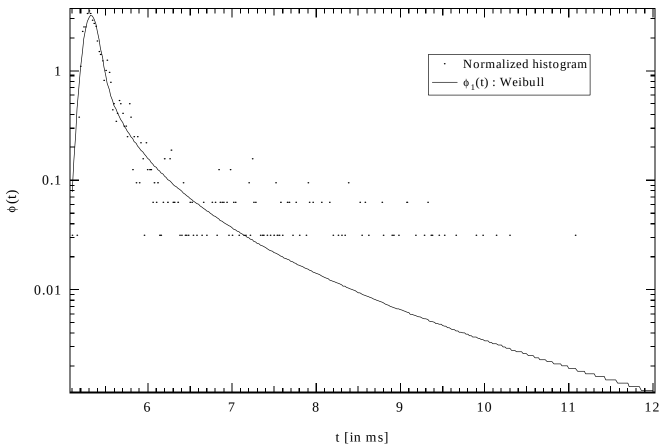

Fig. 10. Both measurement data and the optimal Weibull-fit

C. A Polynomial (Pareto-like) density for $\varphi_1(t)$

We now assume that the entire probability density (and not only the tail) of the closure period behaves polynomially, as a Pareto law, with

$$\varphi_1(t) = \frac{\alpha}{\tau}\left(1+\frac{t}{\tau}\right)^{-\alpha-1}, \qquad t \geq 0, \tag{12}$$

where an additional parameter τ is introduced to fit the Internet traffic for intermediate values of t. Observe that $\int_0^\infty t\varphi_1(t)dt = \frac{\tau}{\alpha-1}$, implying that $\alpha > 1$ in order for the mean to exist. As shown previously, the parameters μ and σ of the normal density φ_2 in (7) are estimated using the night-data. The final three parameters p, τ and α are again estimated by the ML method. For different τ with $\tau = 0.252$ the overall optimum, we find

p	τ	α
0.49	1	2.59
0.45	0.5	1.72
0.38	0.25	1.21
0.39	0.27	1.25

where in the last row $\alpha = 1.25$, the Hill estimate obtained in sec. IV-D.2 below, has been fixed.

These fits are shown in Figure 11.

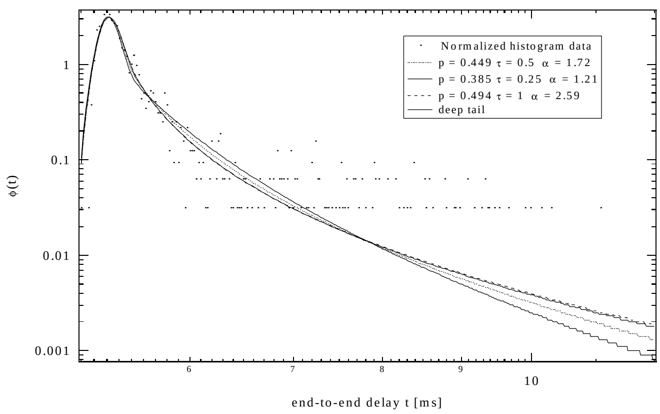

Fig. 11. Both measurement data and the polynomial model on a log-log scale. Clearly, the deep tail coincides with the overall optimum.

D. The tail analysis

In this subsection we apply three different methods to investigate the tail-behavior of our delay measurements: (i) the mean excess, (ii) the Hill estimator and (iii) the α-stable method. We focus on the same data as before, consisting of the $n = 1598$ measurements during the day-period of March 3, 2001, between the boxes A and B.

As before we consider these measurements as realisations $t_1, t_2, \ldots, t_n$ of a statistical sample $T_1, T_2, \ldots, T_n$ from the density (6). The minimum delay will be denoted by m (its realisation equals 5.11 ms.); we will denote the mean of the sample by $\overline{T}_n$, the mean of the data (the realisation of $\overline{T}_n$) is denoted by $\overline{t}_n$ and is equal to 5.61 ms.

Since the Gaussian density φ_2 has a light tail, $\varphi_1 * \varphi_2(t)$ has the same tail behavior as φ_1, if φ_1 exhibits a heavy tail. Specifically (cf. [3]), if

$$\varphi_1(t) \sim c_1 \cdot t^{-(1+\alpha)}, \qquad t \to \infty, \tag{13}$$

then

$$\lim_{t\to\infty} t^{1+\alpha}\varphi_1 * \varphi_2(t) = c_2, \tag{14}$$

for some unspecified constant c_2. We conclude that polynomial tail behavior of the density φ_1, which models the stochastic delay D_T caused by Internet traffic, implies that also the total stochastic delay modeled by (6) has such tail behavior with the same parameter α. Therefore, for a tail analysis, we can concentrate directly on the data $t_1, t_2, \ldots, t_n$.

D.1 The mean excess method

The mean excess function of a random variable X, with right end point $+\infty$ is for $u > 0$ defined by

$$\mathrm{me}(u) = E\left[X - u | X > u\right]$$

For a Pareto random variable with density specified in (12), the function me(u) is linear, with positive slope $\frac{1}{\alpha-1}$ (for $\alpha > 1$). Indeed, if the complementary distribution function of the Pareto is defined as

$$1 - \Phi_1(t) = \left(1 + \frac{t}{\tau}\right)^{-\alpha}, \quad t > 0,$$

then we obtain for a Pareto random variable,

$$\begin{aligned} \mathrm{me}(u) &= \frac{1}{1-\Phi_1(u)} \int_u^\infty (1 - \Phi_1(y))\, dy \\ &= \frac{\tau}{\alpha - 1} + \frac{u}{\alpha - 1}. \end{aligned}$$

One can consult Figure 6.2.4 of Embrechts *et al.* [9] for the form of the mean excess function for other standard distributions.

The mean excess function is a graphical tool to decide between different tail behaviour. In Figure 12, we plot the mean excess for 20 simulations (each of size $n = 1598$) of the Pareto distribution with (the optimal) parameters $\alpha = 1.25$ and $\tau = .27$, together with the *empirical* mean excess obtained from the data. The empirical mean excess function is obtained from the formula

$$\mathrm{me}_n(u) = \frac{1}{1-\Phi_n(u)} \int_u^\infty (1 - \Phi_n(x))\, dx,$$

where Φ_n is the empirical d.f. of $t_1 - m, t_2 - m, \ldots, t_n - m$, the raw data diminished with the minimum delay m. If we fit a straight line in a linear region of the empirical mean excess, that is the region not influenced by the normal distribution (the segment between 1 and 5 ms) as illustrated in the insert in Figure 12, we find $1.55 - 0.23t$, and hence a *negative* slope. It is difficult to draw conclusions from this. As observed from the wide variability in the mean excess of the 20 simulations in Figure 12, we feel that the mean excess plot is doubtful to provide a reliable (or meaningful) estimate of the heavy tail parameter α.

D.2 The Hill estimator.

A reliable method to estimate the tail index in the presence of a heavy tail is the Hill estimator. Suppose again that the tail of the Internet delay density $\varphi_1(t)$ satisfies (13). The extreme value index $\gamma = 1/\alpha$, can be estimated with the Hill estimator [12]:

$$\hat{\gamma} = \left[\frac{1}{k}\sum_{i=1}^{k} \log\left(t_{(n-i+1)} - m\right) - \log\left(t_{(n-k)} - m\right)\right],$$

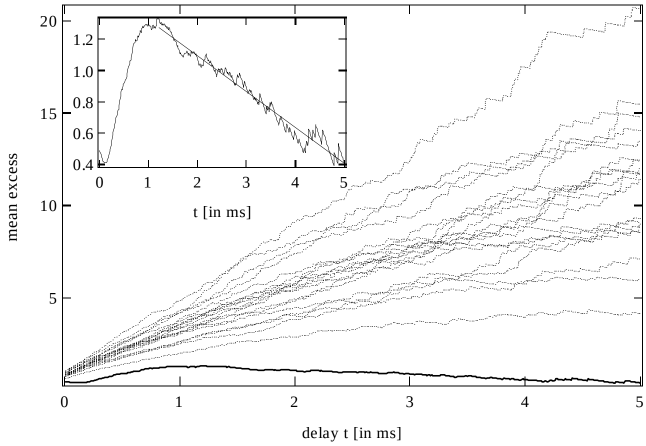

Fig. 12. Mean excess for 20 simulations of a Pareto distribution (12) with $\alpha = 1.25$ and $\tau = 0.27$ in dotted line and the data in bold line. The insert shows mean excess of the data versus delay t. The linear fit satisfies 1.55-0.23t.

where $t_{(1)}, t_{(2)}, \ldots, t_{(n)}$ are the order statistics of the data (the sorted data) and, where k is the number of highest order statistics that are used. If the Hill estimator $\hat{\gamma}(k)$ stabilizes for increasing k to a consistent value $\hat{\gamma}$, as is the case with our data (see Figure 13), this value $\hat{\gamma}$ (here equal to 0.8) provides an estimate for the power $\alpha = \frac{1}{\hat{\gamma}}$, thus here $\widehat{\alpha} = 1.25$. However, as discussed in [6], if the underlying process does not obey a pure power law, the Hill estimator may not be adequate. This is indeed questionable if we inspect the log-log plot of the complementary distribution function as shown in full line in Figure 14. Figure 13 shows

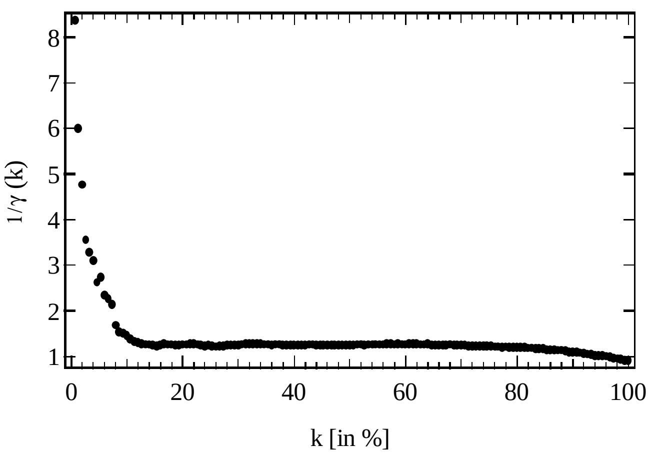

Fig. 13. The behavior of $\gamma(k)$ as function of k where k is expressed in %, i.e. $k = 100\%$ implies that all $n = 1598$ data samples are used.

that over a large range $\hat{\alpha} = 1.25$, which is consistent with the polynomial fit that led to $\alpha = 1.21$, as demonstrated in Section IV-C. Finally, the Hill-estimator $\hat{\gamma} = \hat{\gamma}_{n,k}$ is known to be consistent (meaning that if the sample size n increases the estimator $\hat{\gamma}$ will converge in probability to the true value γ, cf. [15]).

D.3 The $\alpha-$stable method

In a recent paper [6], Crovella and Taqqu describe a procedure to determine the tail index α for distributions of the form

$$\Pr[X > x] = x^{-\alpha} L(x),$$

where L is a slowly varying function at infinity, i.e.,

$$\lim_{t\to\infty} \frac{L(tx)}{L(t)} = 1,$$

for each fixed $x \geq x_0$.

The procedure estimates α by aggregating the centered data $u_1 = t_1 - \bar{t}_n, u_2 = t_2 - \bar{t}_n, \ldots, u_n = t_n - \bar{t}_n$. More precisely it gives estimates for α, based on $u_1^{(k)}, \ldots, u_{n_k}^{(k)}$, for $k = 1, 2^1, \ldots, 2^d$, where

$$u_i^{(k)} = \sum_{j=(i-1)k+1}^{ik} u_j.$$

Figure 14 displays for $d \leq 5$ the tails of the empirical complementary distributions on log-log scale. For an α-stable law these plots should all be parallel with slope $-\alpha$. Comparison with the plots given in Crovella and Taqqu [6] learns that the centered end-to-end delay data $u_1, u_2, \ldots, u_n$ does not satisfy a clean α-stable law, not even in the tail. If one would insist on an estimate of α, curve fitting gives $\alpha = 2.99$, based on $k = 2$ ($d = 1$) and $k = 4$ ($d = 2$), which run more or less parallel. This value is *not* comparable with the value obtained from the Hill estimator $\alpha_{Hill} = 1.25$.

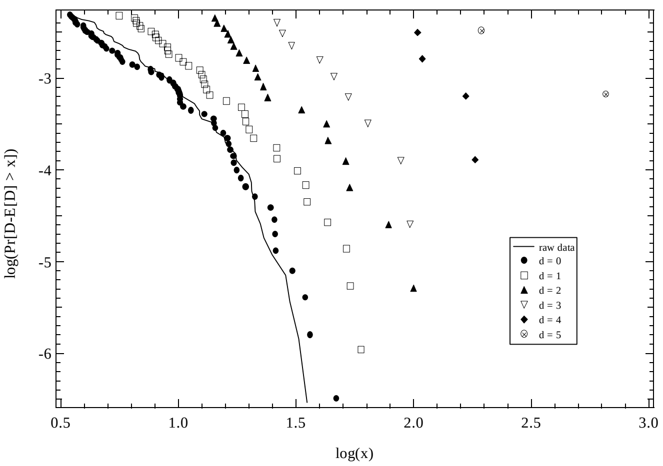

Fig. 14. The α-stable method where 10% of the largest delay data has been used.

E. Non-parametric maximum likelyhood estimation

If we assume that the processing delay φ_2 is specified by a Gaussian probability density (7) with mean μ and variance σ^2, the distribution of the stochastic end-to-end delay D_T caused by Internet traffic can be obtained by deconvolution. We first subtract from the delay sample $T_1, T_2, \ldots, T_n$ the known mean μ to obtain the translated sample

$$Z_1 = T_1 - \mu, Z_2 = T_2 - \mu, \ldots, Z_n = T_n - \mu.$$

According to (6), $Z_1, \ldots, Z_n$ is a sample from the density

$$p g_\sigma(z) + (1-p)(\varphi_1 * g_\sigma)(z), \qquad (15)$$

where $g_\sigma(z) = \frac{\exp(-x^2/2\sigma^2)}{\sigma\sqrt{2\pi}}$ is a Gaussian density with zero mean and variance σ^2. Let us denote the cumulative distribution function (c.d.f.) of D_T by $H(t) = p + (1-p)\int_0^t \varphi_1(u)\,du$, then equation (15) reads (using Lebesgue-Stieltjes integrals)

$$\int g_\sigma(z-x)\,dH(x), \qquad (16)$$

where the unknown H is concentrated on the interval $[0, \infty)$. Groeneboom and Wellner [11, Chapter 2] have presented a method to find a non-parametric maximum likelihood estimator (NPMLE) $\hat{H}_n$ for the unknown c.d.f. H. These authors show that $\hat{H}_n$ satisfies the so-called "self consistency equation" [11, Eq. (2.4)],

$$\hat{H}_n(t) = \frac{1}{n}\sum_{i=1}^{n} \frac{\int_{[0,t]} g_\sigma(z_i - x)\,d\hat{H}_n(x)}{\int_{[0,\infty)} g_\sigma(z_i - x)\,d\hat{H}_n(x)}, \qquad (17)$$

where the sample $Z_1, \ldots, Z_n$ assumes the values $z_1, \ldots z_n$.

To calculate $\hat{H}_n$, one can use the EM (expectation-maximization) algorithm [11, Chapter 3] which is known to converge to the NPMLE [23]. Since our variance σ^2 is small, the EM algorithm converges rather quickly. In our case the EM algorithm is equivalent to iterating the self consistency equation (17),

$$\hat{H}_n^{(k+1)}(t) = \frac{1}{n}\sum_{i=1}^{n} \frac{\int_{[0,t]} g_\sigma(z_i - x)\,d\hat{H}_n^{(k)}(x)}{\int_{[0,\infty)} g_\sigma(z_i - x)\,d\hat{H}_n^{(k)}(x)},$$

where $\hat{H}_n^{(k)}$ is the k-th iterative approximation of $\hat{H}_n$. As a result we obtain the NPMLE $\hat{H}_n$, especially the estimators for the probability p (no Internet traffic) and estimators of the quantiles of c.d.f. H.

From Figure 15 we find estimates for $p = 0.55$, the probability that there is no Internet traffic on the current path and for the quantiles. For instance, with probability 0.95 the end-to-end delay caused by Internet traffic does not exceed 1.92 ms, while with probability 0.99 that delay does not exceed 3.75 ms. The power of this non-parametric method lies in the fact that no model for φ_1 needs to be proposed. Moreover, it is more stable than a deconvolution process based on transforms. Observe, however, that the NPMLE $\hat{H}_n$ has a finite endpoint. More precisely [11, p. 55], $\hat{H}_n(z_{(n)}) = 1$, where $z_{(n)} = t_{(n)} - \mu = D_{\max}$ is the largest data point.

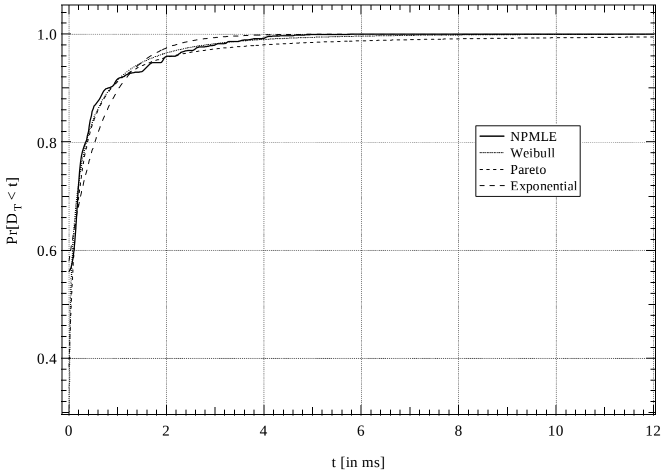

Fig. 15. The c.d.f of the various presented models.

F. Comparison of the different models.

Four different estimates for the delay due to Internet traffic have been investigated,
(i) Exponential, with parameter λ,
(ii) Pareto, with parameters τ and α,
(iii) Weibull, with parameters a and b,
(iv) NPMLE (the non-parametric maximum likelihood estimator).

Figure 15 summarizes the results of these various models. Each of these estimates consists of a density function φ_1 and a probability p. Using our model assumption (6), with the normal density φ_2, the above cases give an estimate of the density φ of the total delay D. For any interval $[s,t]$, an estimate $e_{[s,t]}$ for the expected number of data points of the total n that are contained in the interval $[s,t]$ yields

$$\begin{aligned} e_{[s,t]} &= n\int_s^t \varphi(u)\,du \\ &= np\int_s^t \varphi_2(u)\,du + n(1-p)\int_s^t \varphi_1 * \varphi_2(u)\,du, \end{aligned}$$

This allows us to compare the above established model estimates for the Internet traffic with the RIPE data. For each of the 4 choices (with for each parametric case the optimal parameters), we have calculated the expected frequency e_i of a given cell i, and have compared this frequency with the observed frequency o_i of the data on the fixed path from box A to box B. For these two boxes all data is between 5 and 12 ms and the width of a cell was chosen equal to 0.02 ms. For instance for case (i), we took $\lambda = 1.39$ and $p = .58$ and have calculated for cells c_i in the range from 5 to 12 ms, the expected value as

$$e_i = n\int_{c_i}\left\{p\varphi_2(u) + \lambda(1-p)\int_{-\infty}^{u} e^{-\lambda(u-v)}\varphi_2(v)\,dv\right\}du.$$

Here $n = 1598$, since this is the total number of packets sent during day time, and for the normal density φ_2 we used the parameters $\mu = 5.30$ ms and $\sigma = 0.078$ ms. A statistically sound procedure consists of calculating the χ^2-statistic

$$\chi^2 = \sum_{i=1}^{k}(e_i - o_i)^2/e_i, \tag{18}$$

where k is the total number of cells with positive expected value e_i. Here, the number $k = \frac{12-5}{0.02} + 1 = 351$, because all 4 models are continuous, and hence, all cells between 5 and 12 have positive expectation. The table below presents the results

	Expon.	Pareto	Weibull	NPMLE
χ^2	901	516	531	302

The best fit (lowest χ^2) is reached by the NPML estimator. The expectation of a χ^2-statistic based on k cells equals $k-1$, its standard deviation is given by $\sqrt{2(k-1)}$. The χ^2-statistic with k large can be approximated by a normal distribution, with the same mean and variance, which directly supplies the probabilities for the goodness of the fit. It evidences against the correctness of the parametric models. Only the NPMLE is in accordance with the expected value.

V. Analysis on other fixed paths.

According to the methodology explained in previous sections, the end-to-end delay of other paths has been analysed and the stochastic delay D_T due to Internet traffic has been extracted using the non-parametric method. The results are summarized in the table below.

pathID	h	μ [ms]	σ [ms]	p	95%	99%
1	10	5.30	0.078	0.55	1.92	3.75
2	16	19.45	0.116	0.43	4.98	9.16
3	12	24.97	0.075	0.01	77.0	95.2
4	15	19.09	0.077	0.64	1.65	2.85
5	7	0.98	0.050	0.00	5.20	12.1

The last two columns give the delay value in ms corresponding with 95% and 99% quantile, respectively. The *standard deviation* σ_{proc} of the processing delay per router equals $\sigma_{proc} = \frac{\sigma}{\sqrt{h}}$, which yields for the path under consideration,

pathID	σ_{proc} [μs]
1	25
2	29
3	22
4	20
5	19

The measured σ_{proc} at RIPE NCC equals 10 μs with one router and 14 μs per router, when measured in a configuration with two routers. Papagiannaki *et al.* [19] report closely agreeing values of σ_{proc} ranging from $13\mu s$ to $24\mu s$.

Figure 16 compares the distribution $\Pr[D_T < xD_{\mathsf{max}}]$ (where $x \in [0,1]$) for these paths and illustrates a large variety in the behavior. For small values of x, the probability p that the probe-packet is not delayed is decisive. Small values of p may hint towards fairly loaded paths, while high values of p towards non-congested or over-dimensioned paths. Figure 17 shows in essence the same information as in Figure 16 on different scales. Especially the intermediate region of $\Pr[D_T > xD_{\mathsf{max}}] = 1-$

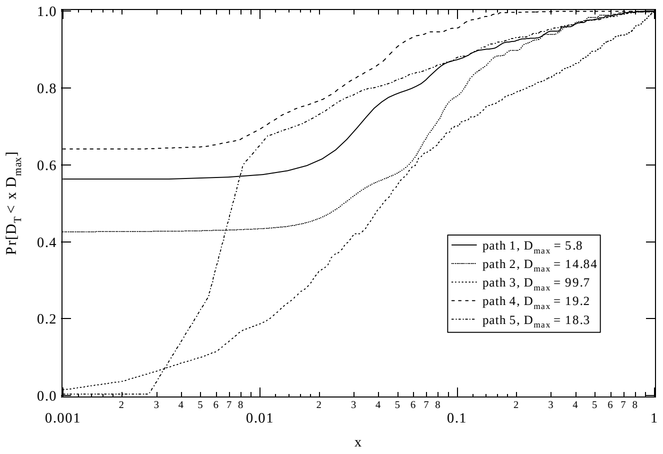

Fig. 16. The c.d.f of the random delay due to Internet traffic for 5 different paths versus the normalized delay $x = \frac{d}{D_{\mathsf{max}}}$.

$\Pr[D_T < xD_{\mathsf{max}}]$ is observed to be reasonably linear. After fitting in the largest possible region where the function $\ln(\Pr[D_T > xD_{\mathsf{max}}]) = a_1 x + a_2$ is appearently linear, we found for the different paths the following results

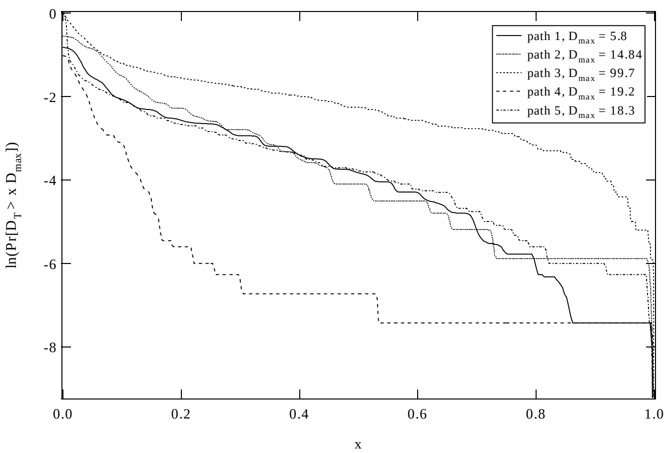

Fig. 17. The logarithm of $\Pr[D_T > xD_{\mathsf{max}}]$ versus x for the five paths.

pathID	a_1	a_2
1	-5.49	-1.3
2	-5.93	-1.1
3	-2.53	-1.0
4	-22.8	-1.0
5	-4.51	-1.6

Clearly, $\Pr[D_T > xD_{\mathsf{max}}]$ for paths 1, 2 and 5 behave fairly well exponentially with rate around -5. The exponential behavior of path 3 and 4 seems less pronounced. Finally, the deep tail is not exponential for any of these paths. Only the tail of path 4 was sufficiently linear on $\ln(\Pr[D_T > xD_{\mathsf{max}}])$ versus $\ln x$ plot with slope -1.34. The other paths were not convincingly linear on that plot. In summary, there does not seem to exist a single, simple shape for the delay caused by Internet traffic on fixed paths.

We may conclude that either much more measurements (than about 2150 in our case) are needed to determine the tail region or that the end-to-end delay measurements on a fixed path are not the best way to detect the underlying physical (long range or self-similar) mechanisms.

VI. Conclusions

The stochastic end-to-end delay of IP probe-packets on fixed paths has been investigated between various measurement boxes scattered over mainly a European part of the Internet. That stochastic end-to-end delay was found to consist of two components: (i) a router processing delay and (ii) queueing delay D_T caused by Internet traffic. From lab-measurements presented elsewhere [5], we succeeded in identifying the distribution of the router processing delay as the independent sum of symmetric shaped delay variables, which are satisfactorily approximated by a Gaussian distribution.

The queueing delay D_T due to Internet traffic was modeled by on/off processes to explain the blocking of the IP probe-packet over the fixed path. The combination of both router processing delay and queueing delay lead to the basic model (6) for the total delay.

For the queueing delay D_T, represented by $\varphi_1(t)$ in the basic model (6), we have proposed to examine the exponential, the Pareto and the Weibull distribution. Using the maximum likelihood method, best fit parameters were obtained that, unfortunately, seem to lead to poor representation of the orginal data. One explanation may be sought in the limited number of intermediate or 'tail' observations. Another suggestion is that neither of the proposed models seems to capture the specific tail behavior accurately enough. At last, the analysis has implicitly assumed that the Internet traffic was stationary (or in steady state), characterized by one distribution function $\varphi_1(t)$ during the measurement period.

On the other hand, we succeeded in deconvolving the stochastic end-to-end delay to obtain a non-parametric maximum likelihood estimate (NPMLE) of the distribution function of the delay D_T caused by Internet traffic. The interest of the NPMLE lies in the accurate estimate of the probability p that a probe-packet is *not* delayed by Internet traffic. This performance measure p of the path between the two boxes also reflects the probability that the path is unloaded. The quantiles of the NPMLE give pre-designed probabilities bounds for the queueing delay caused by Internet traffic. These bounds are valuable to deploy quality real-time services on Internet, in particular the well-studied and basic telephony over Internet service.

Acknowledgements. We thank Dr. Henk Uijterwaal of RIPE NCC for the data and his support with the measurements of the router processing delay. We thank Harry Mertodimedjo and Charles Bovy, two of our students, for their help in analyzing the data.

References

[1] M. Abramowitz and I.A. Stegun, Handbook of Mathematical Functions, Dover, 1968.

[2] A. Adams, T. Bu, T. Friedman, J. Horowitz, D. Towsley, R. Caceres, N. Duffield, F.L. Presti, S. B. Moon, V. Paxson, "The use of end-to end multicast measurements for characterizing internal network behavior", IEEE Communications magazine, vol. 38, 5, pp. 152–159, 2000.

[3] N.H. Bingham, C.M. Goldie and J.L. Teugels, Regular Variation, Cambridge University Press, Cambridge, 1987.

[4] J.C. Bolot, "Characterizing end-to-end packet delay and loss in the Internet", J. High-Speed Networks, vol. 2, no 3, pp. 289-298, 1993.

[5] C. J. Bovy, H. T. Mertodimedjo, G. Hooghiemstra, H. Uijterwaal and P. Van Mieghem, "Analysis of end-to-end delay measurements in Internet", submitted to PAM 2002.

[6] M. E. Crovella and M. S. Taqqu, "Estimating the heavy tail index from scaling properties", Methodology and Computing in Applied Probability, vol. 1, no. 1, pp. 1-21, 1999.

[7] D. De Vleeschauwer, A. Van Moffaert, J. Janssen, M. J. C. Büchli, G. H. Petit, B. Steyaert and H. Bruneel, "Determining the number of packet-based phones that can be supported by one access node", Proceedings of the 14th ITC Specialists Seminar on Access Networks and Systems (ITCSS01), Girona (Spain), 25-27 April, pp. 197-204, 2001.

[8] N.G. Duffield and P.F. Lo Presto, "Multicast inference of packet delay variance at interior network links", IEEE Infocom 2000, Tel Aviv, Israel, 2000.

[9] P. Embrechts, C. Klüppelber and T. Mikosch, Modelling Extremal Events, Springer Verlag, 1997.

[10] A. Fei, G. Pei, R. Liu and L. Zhang, "Measurements on delay and hop-count of the Internet", Proc. of the IEEE Global Internet '98, December 1998.

[11] P. Groeneboom and J.A. Wellner, Information Bounds and Nonparametric Maximum Likelihood Estimation, Birkhäuser 1992, Basel.

[12] B.M. Hill, "A simple general approach to inference about the tail of a distribution", Annals of Statistics, vol. 3, pp 1163–1174, 1975.

[13] R. Caceres, N.G. Duffield, J. Horowitz and D. Towsley, "Multicast-based inference of network internal loss characteristics", To appear in IEEE TON.

[14] M. Falk, J. Hüsler and R.D. Reiss, Laws of Small Numbers: Extremes and Rare events, Birkhauser, Basel, 1994.

[15] D.M. Mason, "Laws of large numbers for sums of extreme values", Annals of Probab., vol. 1, pp 754–764, 1982.

[16] V. Paxson, Measurements and Analysis of end-to-end Internet dynamics, Ph.D. Dissertation, University of California at Berkeley.

[17] F. Georgatos, F. Gruber, D. Karrenberg, M. Santcroos, A. Susanj, H. Uijterwaal and R. Wilhem, "Providing Active Measurements as a Regular Service for ISP's", Proceedings of Passive and Active Measurement (PAM2001), Amsterdam, The Netherlands, April 23-24, pp. 45-56, 2001.

[18] F. Poppe, D. De Vleeschauwer and G. H. Petit, "Choosing the UMTS Air Interface Parameters, the Voice Packet Size and the Dejittering Delay for a Voice-over-IP Call between a UMTS and a PSTN party", IEEE INFOCOM01, Anchorage (USA), April 22-26, pp. 805-814, 2001.

[19] K. Papagiannaki, S. Moon, C. Fraleigh, P. Thiran, F. Tobagi and C. Diot, "Analysis of Measured Single-Hop Delay from an Operational Backbone Network", submitted to IEEE INFOCOM2002.

[20] P. Van Mieghem, "A Lower Bound for the End-to-End Delay in Networks: Application to voice over IP", Globecom'98, Nov. 8-12, Sydney (Australia), pp. 2508-2513, 1998.

[21] A. Van Moffaert, D. De Vleeschauwer, J. Janssen, M. J. C. Büchli, G. H. Petit and P. Coppens, "Tuning the VoIP Gateways to Transport International Voice Calls over a Best-Effort IP Backbone", Proceedings of the 9th IFIP Conference on Performance Modelling and Evaluation of ATM & IP Networks (IFIP01), Budapest (Hungary), 27-29 June, pp. 193-205, 2001.

[22] W. Willinger, M. S. Taqqu, R. Sherman, R. and D. V. Wilson, "Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level", IEEE/ACM Transactions on Networking, vol. 5, No. 1, pp. 71-86, Feb. 1997.

[23] C.F.J. Wu, "On the convergence properties of the EM algorithm", Annals of Statist., vol. 11, pp 95–103, 1983.

VII. Appendix

Proof of Theorem 1: Consider the case where $X_1 = x$ and $Y_1 = y$ are given. If $X_1 + Y_1 = x + y$ happens before time t, so if $x + y < t$, then the first renewal of occurs before time t and the event $\{\beta_t > u | X_1 + Y_1 = x + y\}$ has the same probability as the event $\{\beta_{t-x-y} > u\}$, because the renewal process starts from scratch at time $x + y$. If on the other hand $x + y > t$, the residual closure time β_t is contained in the first renewal interval $[0, x + y]$. In this case we can only have $\beta_t > u$ if $t > x$ and $x + y - t > u$ hold simultaneously. Hence,

$$\Pr[\beta_t > u | X_1 = x, Y_1 = y] = \Pr\left[\beta_{t-x-y} > u\right],$$

if $x + y < t$ and

$$\Pr[\beta_t > u | X_1 = x, Y_1 = y] = I(t > x) \cdot I(x + y - t > u),$$

if $x + y > t$.

Denoting for u fixed, $Z(t) = \Pr[\beta_t > u]$ we obtain:

$$\begin{aligned} Z(t) &= \int_x \int_y \Pr[\beta_t > u | X_1 = x, Y_1 = y] f(x)g(y)\,dxdy \\ &= \int \int_{x+y<t} \Pr\left[\beta_{t-x-y} > u\right] f(x)g(y)\,dxdy \\ &\quad + \int_{x<t} \int_{y>t+u-x} f(x)g(y)\,dydx \end{aligned}$$

or equivalently, with the notation:

$$\begin{aligned} C(t) &= F * G(t) = \int_0^t F(t-s)dG(s) \\ &= \int_0^t F(t-s)g(s)\,ds, \\ Z * C(t) &= \int_0^t Z(t-s)dC(s) \\ &= \int_0^t Z(t-s)[\int_0^s f(s-v)g(v)\,dv]\,ds, \end{aligned}$$

the renewal equation:

$$Z(t) = z(t) + Z * (F * G)(t),$$

where

$$\begin{aligned} z(t) &= \Pr[X < t, Y > t + u - x] \\ &= \int_{x<t} \int_{y>t+u-x} f(x)g(y)\,dydx. \end{aligned}$$

By the key-renewal theorem:

$$\begin{aligned} \lim_{t\to\infty} \Pr[\beta_t > u] &= \frac{1}{\mu_F + \mu_G} \int_0^\infty z(t)\,dt \\ &= \frac{1}{\mu_F + \mu_G} \int_0^\infty \int_0^t f(s)[1 - G(u+t-s)]\,ds\,dt \\ &= \frac{1}{\mu_F + \mu_G} \int_0^\infty f(s) \int_s^\infty [1 - G(u+t-s)]\,dt\,ds \\ &= \frac{1}{\mu_F + \mu_G} \int_0^\infty f(s) \int_0^\infty [1 - G(u+y)]\,dy\,ds \\ &= \frac{1}{\mu_F + \mu_G} \int_u^\infty [1 - G(y)]\,dy. \end{aligned}$$

This proves Theorem 1. ¤