Visualization Challenges in Internet Traffic Research

Bárbara González–Arévalo
Department of Mathematics
University of Louisiana at Lafayette
Lafayette, LA 70504-1010

Félix Hernández-Campos
Department of Computer Science
University of North Carolina
Chapel Hill, NC 27599-3175

J. S. Marron
Department of Statistics and Operations Research
University of North Carolina
Chapel Hill, NC 27599-3260

Cheolwoo Park
Department of Statistics
University of Florida
Gainesville, FL 32611-8545

September 19, 2004

Abstract

This is an overview of some recent research, and of some open problems, in the visualization of internet traffic data. One challenge comes from the sheer scale of the data, where millions (and far more if desired) of observations are frequently available. Another challenge comes from ubiquitous heavy tail distributions, which render standard ideas such as "random sampling will give a representative sample" obsolete. Some alternate sampling approaches are suggested and studied. One more challenge is the visual representation of (and even the definition of) "common constant transfer rates" in a large scatterplot.

1 Introduction and background

The area of Internet traffic measurement and modelling has a pressing need for novel and creative visualization ideas. The issues and the data are both complex, yet few researchers in that area (with some notable exceptions) are aware of the power of visualization for addressing the problems, and understanding complicated behavior.

The Internet shares some similarities to the telephone network. Both are gigantic, worldwide networks for the transmission of information. Both share the notion of "connection", generally between two points. For this reason the first models for Internet traffic were based on standard queueing theory, with assumption of Poisson arrival of connections, and of exponentially distributed times of connection duration.

A large body of exciting work during the 1990's revealed that these assumptions were grossly inadequate, and far different models were usually much more appropriate. In particular, duration distributions were seen to exhibit heavy tails (caused by both far shorter, and also far longer connections than typically found in telephone traffic), and time series of aggregated traffic exhibit bursty behavior and long range dependence. An elegant mathematical theory, demonstrating how heavy tail durations can lead to long range dependence was developed by Mandelbrot (1969), Cox (1984), Taqqu and Levy (1986), Leland, Taqqu, Willinger and Wilson (1994), Crovella and Bestavros (1996), Heath, Resnick and Samorodnitsky (1998), and Resnick and Samorodnitsky (1999).

This theory is deep and compelling, and gives good description of observed behavior in a wide range of circumstances. However, there has been recent controversy at several points.

One controversial point has been the issue of the heaviness of tails. Downey (2000) suggests that the Log Normal (not heavy tailed in the classical sense), can fit duration distributions as well as classical heavy tail distributions, and gave some interesting physical motivation for this distribution as well. However, by developing the nice idea of *tail fragility*, Gong, Liu, Misra and Towsley (2001) showed that both types of distribution can give apparently reasonable fits. This general direction was further developed by Hernández-Campos, Marron, Samorodnitsky and Smith (2002), using some much larger data sets (in the millions), together with a novel visualization for understanding the level of sample variation. This latter work showed that a mixture of three Double Pareto Log Normal distributions (see Reed and Jorgensen 2004) gave an excellent fit, and was also physically interpretable. These results motivated the development of the concept of *variable heavy tails*.

Another point of recent controversy has been over the issue of long range dependence. This is currently widely accepted (and intuitively sensible), but some interesting questions have been raised (using some novel visualization ideas) by Cao, Cleveland, Lin, and Sun (2001, 2002a, b, c). The key idea is that aggregated traffic, of the type typically found at major Internet nodes, tends to wash out long range dependence. The idea is theoretically justified by appealing to limit theorems for aggregated point processes, and supported by recent measurement studies that examined high-capacity backbone links, Zhang et al. (2003) and Karagiannis et al. (2004). An example, where both types of behavior were observed, depending on scale, was studied using some different visualizations, by Hannig, Marron and Riedi (2001). An interesting issue to follow in the future will be the state of this balance between long range dependence caused by relatively few but extremely large transmissions, and a more Poisson type probability structure caused by aggregation. Cao, Cleveland, Lin, and Sun (2001, 2002a, b, c) predict ultimate Poisson type structure, for the good reason that Internet traffic continually increases. However, this is based on an assumption that the distribution of sizes of transmissions will stay fixed, which seems questionable.

Downey (2001) questioned long range dependence from a different viewpoint, by showing that duration distributions may not be very consistent with the definition of *heavy tailed*, in the classical asymptotic sense. This was the first observation of *variable heavy tails*, as defined in Hernández-Campos, Marron, Samorodnitsky and Smith (2002). That paper goes on to develop an asymptotic theory that parallels the classical theory. In particular it is seen that long range dependence still follows from the far broader (and realistic in terms of the nature of the data) concept of variable heavy tails.

This paper points out some perhaps fun and challenging visualization problems, and offers some approaches to their solution.

The first main problem is related to the *Mice and Elephants* graphic, developed in Marron, Hernández-Campos and Smith (2002), discussed in Section 2. The problem is how to choose a *representative sample*, and it is seen that the usual device of random sampling is clearly inappropriate. In Section 3, two different biased sampling approaches to this problem are proposed. A quite different approach to visualizing this large and complex population, based on studying quantile windows, is shown to reveal other interesting structure in the data, in Section 4. Our work on flow sampling for visualization complements existing ideas developed in the context of sampled traffic monitoring. Recent work in this area, Duffield and Lund (2003) and Hohn and Veitch (2003), has shown that flow sampling provides a more accurate view of Internet traffic than packet sampling. In particular, Hohn and Veitch have shown that flow sampling preserves long range dependence.

The second main problem is motivated by an apparent "commonality of flow rates", discussed in Section 5. A large scatterplot seems to reveal some interesting visual structure, that makes physical sense. The question is how to best understand the driving phenomena.

2 Mice and Elephant plots and random sampling

The Mice and Elephants plot is a visualization that illustrates the fundamental theory discussed in Section 1. In particular, it shows how a heavy tailed distribution can lead to long range dependence, as explained below. This type of plot is shown in Figure 1. The key idea is that Internet *flows*, i.e. the set of packets that make up a single connection, are represented by line segments. The left (right, resp.) ends of the line segments show the times of the first (last, resp.)

packets in each flow. Thus each line segment represents the *overall time of activity* of that flow. For good visual separation of the line segments, a random height is used on the vertical axis (essentially the *jitter* idea of Tukey and Tukey (1990), see also Cleveland (1993)). This is in contrast with the common packet arrival time vs. address plots commonly used for traffic anomaly detection, see for example Cho (2001).

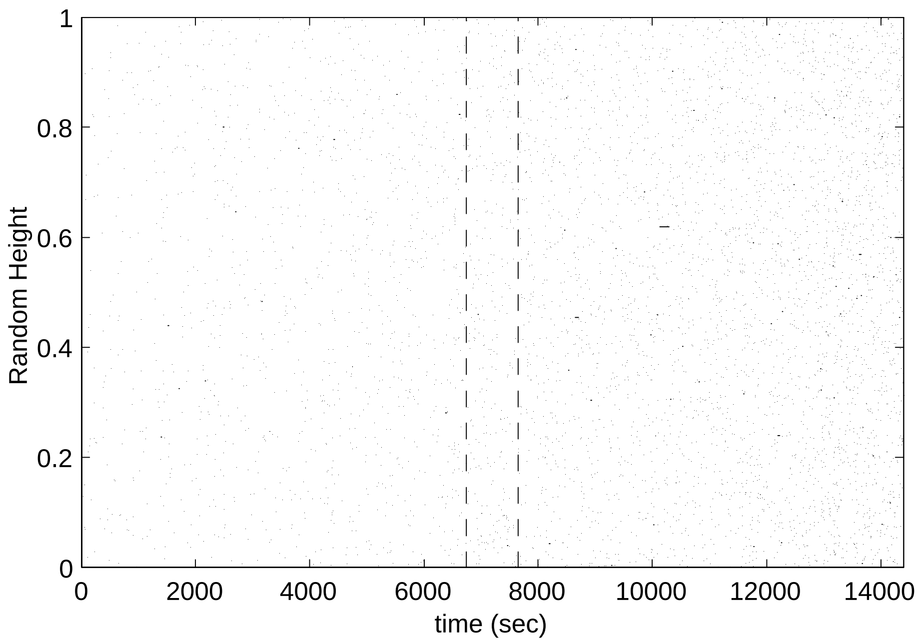

FIGURE 1: *Mice and Elephants plot for full four hour time block. Vertical bars indicate 15 minute time window shown in Figure 2. This suggests that all flows are "mice".*

The data here were HTTP (Web browsing) response times. They were collected during a four hour period, 8:00 AM - 12:00 noon, on a Sunday morning in April of 2001. This time period was chosen to represent a *light traffic* time. For a parallel analysis of a *heavy traffic* time, see Marron, Hernández-Campos, and Smith (2002). More detailed graphics for both analyses are available at the web address: http://www.cs.unc.edu/Research/dirt/proj/marron/MiceElephants/. For more details on the data collection and processing methods, see Smith, Hernández-Campos, Jeffay and Ott (2001).

The total number of flows for the time period in Figure 1 is 1,070,545. Massive overplotting resulted from an attempt to plot all of them. A simple and natural approach to the overplotting problem is to plot only a random subsample. This was done for a subsample size of 5000 (chosen for good visual effect) in Figure 1, and in the other figures in this section.

Figure 1 shows steadily increasing traffic, which is expected behavior on Sunday mornings (perhaps the times at which students begin web browsing is driven by a wide range of adventures experienced on the previous night!). It also suggests that there are no long flows, with the longest visible flow being less than 5 minutes. This is a very serious mis-impression, that completely obscures the most important property of the traffic, as noted below.

This point becomes clear from a similar graphic, but zoomed into the region between vertical bars in Figure 1, which represent the central 15 minutes (1/16th of the total time). Figure 2 shows this zoomed mice and elephants plot. There were 59,113 (not far from 1,070,545 / 16) flows that intersected this time range. Plotting all would again result in severe overplotting, so only a random sample of 5000 is plotted.

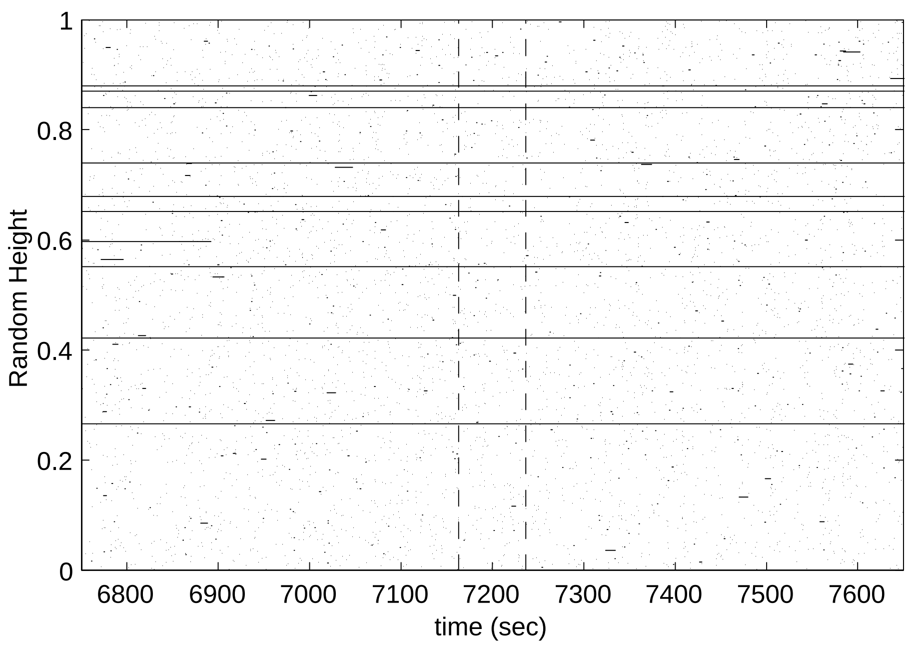

FIGURE 2: *Mice and Elephants plot for 15 minute time block. Vertical bars indicate time span containing 5000 flows, shown in Figure 3. This shows both "mice" and "elephants".*

The visual impression of Figure 2 is far different from that of Figure 1. In particular, there are a number of flows that cross the full 15 minute interval, which seems quite contrary to the visual impression from Figure 1 that all flows are much less than 5 minutes in duration. This mis-impression is caused by a combination of the heavy tailed duration distribution and the random sampling process. Because of the heavy tails, there are only a very few flows that are very long. These have only a very small chance of appearing in the randomly selected sample. E.g. the chance that any of the largest 40 flows have a chance of appearing is only about $(40 \cdot 5000/1,070,545) \approx 0.19$. The number 40 is relevant, since 38 flows extend the full length of the central one hour time interval. This small probability of inclusion explains why none of these very long flows appear in Figure 1.

It is interesting to zoom in once again. Figure 3 shows the results of repeating the visualization for the region between the vertical bars in Figure 2. Those bars do not show 1/16th of the region in Figure 2, because that contains less than 5000 flows (which would give an inconsistent visual representation). Instead the bars are chosen so that exactly 5000 flows intersect the time interval (which is again centered in the range of the data), which is about 1.3 minutes long.

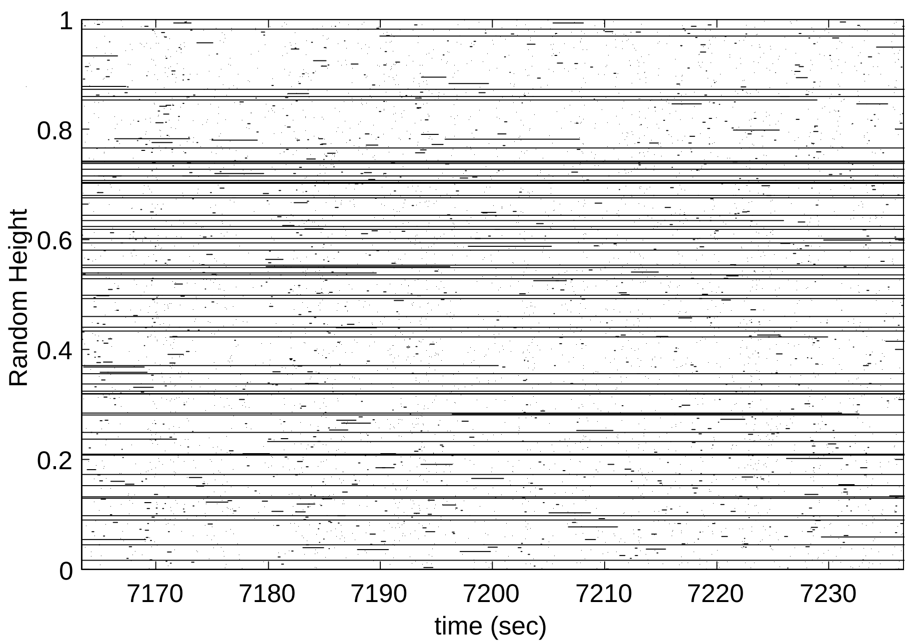

FIGURE 3: *Mice and Elephants plot for time window containing 5000 flows.*

Figure 3 shows a rather large number of long flows, and because there is no sampling, is perhaps *representative of behavior at a given time*. However, this view is also biased because it shows in some sense *too high a proportion of long flows*. The reason is a *length biasing* type of effect: long flows have a much greater chance of appearing in any such small interval, yet as noted above they are a very small fraction of the population.

The clear conclusion from Figures 1, 2, and 3 stands in stark contrast to one of the most time honored principles of statistics (and a commonly used tool in visualization): simple random sampling of these data does not give a *representative sample*. Again this problem is caused by the heavy tails of the duration distribution of Internet connections, and there is a general principle at work: simple random sampling will never give a representative sample in heavy tail situations.

The first open problem proposed in this paper is to find an improved version of *representative sample*. A sensible first step may be to decide what that means. Is there a reasonable mathematical definition of this that makes sense for heavy tailed distributions? Can classical length biased sampling ideas perhaps be useful? Two biased sampling approaches to these problems are given in Section 3. The first approach essentially considers a wide range of zooming window views of the data, of which Figures 1, 2 and 3 are example, and uses probabilities of appearance in the window to assign sampling weights, as discussed in Section 3.1. The second approach assigns weights using the Box Cox transformation of the data, with details appearing in Section 3.2.

Figures 1, 2 and 3 show that the name *Mice and Elephants* is sensible for this graphic. It has become commonplace terminology in the Internet research community for this phenomenon of a very few, very large flows. This concept is fundamental to the ideas outlined in Section 1. It is a clear consequence of the heavy tail duration distributions. It also makes the long range dependence in the aggregated time series visually clear. In particular, time series of binned traffic measures (such as packet counts) are essentially vertical sums of the line segments in the mice and elephants plots. The above described theory about heavy tails implying long range dependence is visually clear, given that the very long elephant flows clearly persist over quite long time ranges. It is not surprising that the persistence of the elephants results in the often observed *bursty behavior* of Internet traffic.

Another interesting open problem is to use this visualization to motivate new quantitative measures for understanding the nature of this type of data. The standard notions of heavy tails for the duration distributions, and of time series dependence, are not aimed at describing the full structure of this data. Instead they are just tools adapted from other areas, which perhaps result in a somewhat clumsy statistical analysis. Can the quantitative analysis be sharpened by quantitating other aspects of the full plot?

Mice and elephants visualizations also give a very clear view of the fact that the standard queueing theory models, with exponential duration distributions, are grossly inappropriate. This is seen in Figure 4, which is a duplicate of Figure 3, but for simulated data, with exponential distributions. To keep the comparison as fair as possible, the real data time range, sample sizes and even start times are used. Only the duration of each flow (the length of the line segment) is simulated. Also the exponential parameter is chosen to give a population mean that is the same as the sample mean for the real data.

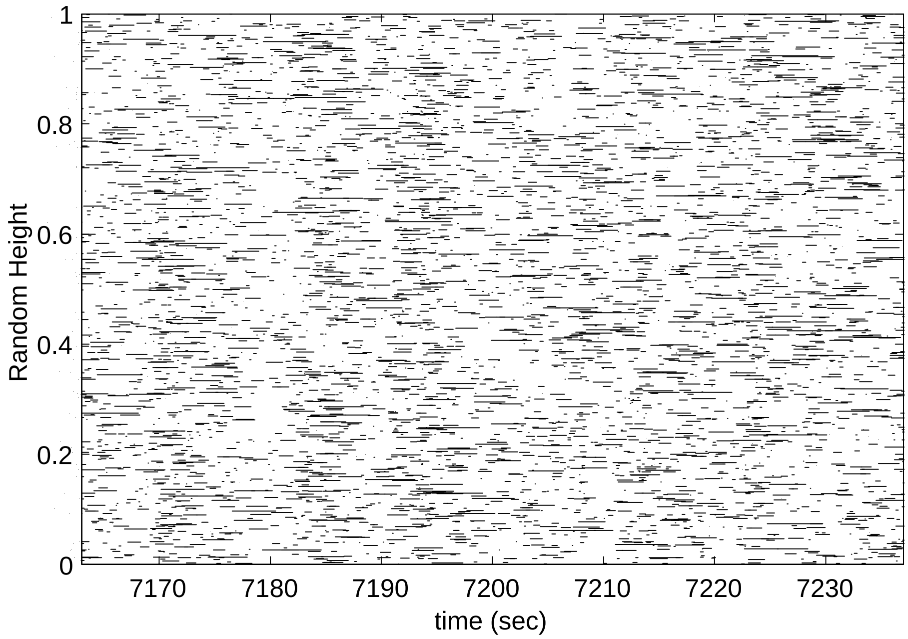

FIGURE 4: *Mice and Elephants plot for simulated exponentials, in setting of Figure 3.*

Figure 4 shows a completely different type of distribution of flow lengths, from the real data shown in Figure 3. In particular, there are no flows that are nearly long enough cover the whole interval, the number of very short flows is far fewer, and there are many more "medium size" flows. This is a consequence of the "light tail" property of the exponential distribution. Once the mean is specified, there are constraints on the frequencies of very large and very small observations. These constraints make the exponential distribution a very poor approximation to the type of behavior seen in Figures 1, 2 and 3. Thus these mice and elephants plots clearly illustrate the concept from Figure 1 that classical queueing models are inappropriate for Internet traffic. In addition, the mice and elephants plot in Figure 4 seems quite consistent with the idea that when this traffic is vertically aggregated, the resulting time series exhibit only classical short range dependence.

The above proposed problem of how to subsample for effective mice and elephants visualization should not be regarded as "one off". The reason is that the Internet is constantly changing in many ways, and this could become a standard tool for monitoring change. For example, such monitoring could show the ultimate resolution to the above controversy, as to whether large scale aggregation will eventually swamp out long range dependence effects, or whether the latter will continue through the continued growth of elephants in frequency and size. An effective solution might also extend well beyond Internet traffic, and might provide the beginnings of a new theory of sampling in heavy tail contexts.

3 Biased Sampling

In this section, two biased sampling methods are proposed, which give samples that give a better visual impression of the nature of this complex population. The first is based on *windowed biased sampling.* The second takes a quite different *Box-Cox biased sampling* approach. Both approaches have connections to classical length biased ideas, but are adapted to the present special setting.

Both approaches provide a visualization over the full time interval, as shown in Figure 1, involving a non-uniform random sample of the full set of flows. But each involves construction of a weighting scheme, that gives the elephants an appropriate probability of appearing.

In this section, we only consider responses with nonzero duration. The number of these responses is $N = 382127$.

3.1 Windowed Biased Sampling

In this section probability weights are chosen to mimic the chance of appearing in a window that is *the right size* among a succession of zooming windows. Figures 1, 2 and 3 are members of this succession of zooming windows.

Figure 5 shows the result of the biased sampling method developed in this section. This shows nearly 5000 (actually 4834) HTTP responses, which are drawn in a way that gives a much clearer impression of the Mice and Elephants nature of this data set.

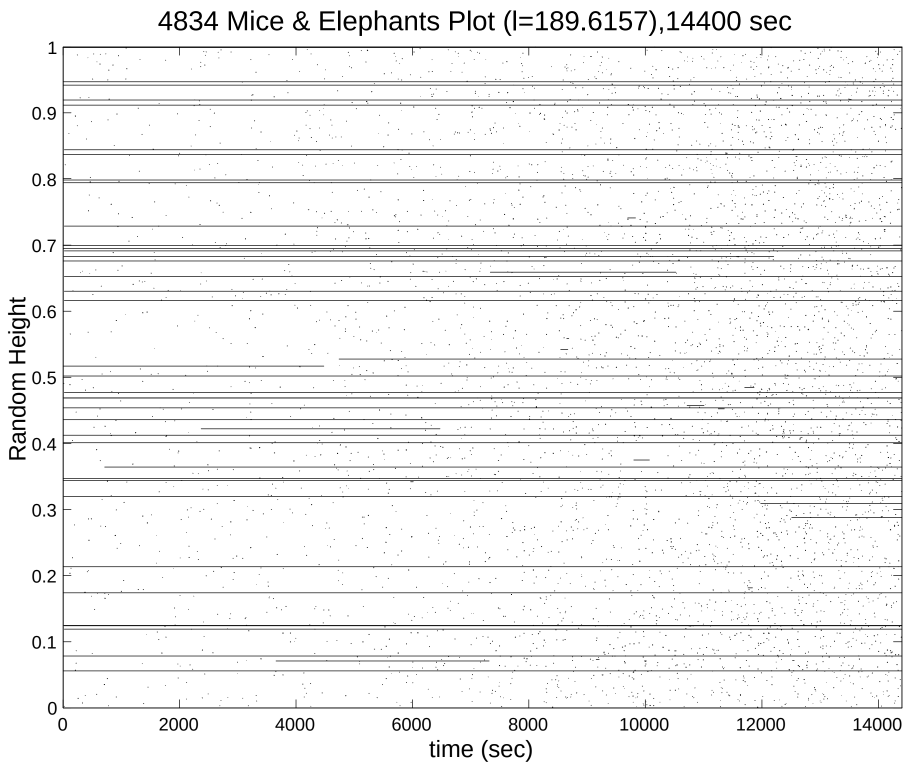

FIGURE 5: *Windowed Biased Sampling of Sunday morning HTTP Response data. Shows both mice and elephants.*

Notation for making this precise, which allows straightforward generalization to other contexts is now introduced. Let L denote the total time period (4 hours for these data, as seen in Figure 1) , and let x and d be the starting time and duration of each HTTP response, i.e. x is the left endpoint and $x+d$ is the right endpoint of each line segment in Figures 1, 2 and 3. Let N denote the total number of HTTP responses. Then the *density of the data* is $\delta = \frac{N}{L}$. Figure 1 shows that the density of the data is clearly not constant over time (there is a clear increasing trend), so δ should really be viewed as the *average density*.

Let n denote the number of responses desired to appear in the final full time scale visualization (taken to be 5000 in Figures 1, 2 and 3). Let l be a candidate subwindow size (i.e. l indexes a zooming family of windows). An interesting view over a wide range of values of l is available in the Windowed Biased Sampling section of the web page González-Arévalo, et al (2004).

An appropriate choice of l will give a good subwindow approximation of the overall average data density δ. Thus it makes sense to use $l = \frac{n}{\delta}$.

The HTTP responses, that appear in the final full time span visualization, shown in Figure 5, are chosen with a weighted random sampling scheme, where the weights are chosen to reflect the probability that the given flow appears in a window of length l.

To calculate these weights, consider a point U chosen at random from the interval $[0, L-l]$. Then we could take all the responses in the interval $[U, U+l]$ as our sample. In this case the sample size will be close to n. But this sample has the drawback that it does not cover the full time range $[0, L]$. To cover the full range, we devise a weighted random sampling scheme as follows.

Because some responses begin before time 0, or finish after time L, we recommend weights that depend on whether or not the response is *fully captured*, meaning that all data packets were transferred, within the interval $[0, L]$. HTTP responses that had some of their data transferred either before the starting time, or after the ending time, are called *censored*.

First we assign weights to the fully captured responses. We model such a response as having a random start time X, assumed to be uniformly distributed on $[0, L]$. Let $p(d, l)$ be the probability that an observation, with duration d, intersects our random window $[U, U+l]$. Let $g(x, d, l)$ be the conditional probability that we choose an observation given that its starting time is x, its duration is d, for the window width l. Let $f(x) = \frac{1}{L}$ be the probability density function of X. Then

$$\begin{aligned} g(x,d,l) &= Pr(U \in [x-l, x+d]) \\ &= \frac{\min(x+d, L-l) - \max(x-l, 0)}{L-l}. \end{aligned} \tag{1}$$

Averaging this over x values, i.e. taking expectation over X, gives an approximate probability that a randomly chosen response intersects the window $[U, U+l]$,

$$\begin{aligned} p(d,l) &= Eg\,(X,d,l) = \int_0^L g(x,d,l) f(x) dx \\ &= 1 - \frac{1}{2}\frac{L-l}{L}\left[1 + \left(1 - \frac{d}{L-l}\right)^2\right], \end{aligned} \tag{2}$$

which we use as the weight for random selection of responses.

Second we assign weights to the censored responses. A similar calculation gives

$$p(d,l) = \min\left\{\frac{d}{L-l}, 1\right\}. \tag{3}$$

Note that any flow with duration at least $L-l$ will have probability one of being in the sample. Thus, all 38 HTTP responses that cover essentially the full 4 hour time interval appear in Figure 5.

So, for fixed l and L we can choose each observation with probability $p(d, l)$ given in (2) and (3). This will give us a sample that spans the whole time interval, but that gives a higher probability to the elephants. A possible drawback of this method is that the size of our sample is random, and depends on l, although it will be close to $n = l\delta$.

When a sample of exactly size n is required, we can do the following:

1. Start with a larger window, with width $l = \frac{n}{0.8 \times \delta}$.

2. Choose each observation with probability $p\,(d, l)$.

3. From this larger sample, choose a random sub-sample of size n.

Note that the number 0.8 is arbitrary, and intended to guarantee that there are at least n points in the sample. However, care should be taken, in the choice of this scale factor, to ensure that the original sampling scheme produces a sample not much larger than the desired n. Otherwise, the subsequent random sampling will tend to reproduce the effect we are trying to avoid: random sampling downweights the elephants.

3.2 Box-Cox Biased Sampling

As discussed in Section 2, random sampling does not provide a useful visualization of durations of HTTP responses because of the heavy-tail of the duration distribution. Here we study a biased sampling approach based upon transformations. In this section, the censored observations are handled differently, this time by simply ignoring whether or not responses extend beyond the range $[0, L]$.

The starting point of this approach to bias sampling, is to represent simple random sampling in terms of the Complementary Cumulative Distribution Function (CCDF), $\bar{F}(d) = P\{D > d\}$ of the duration D. Simulated values can be drawn, by generating Uniform (0,1) realizations, U, and then taking $D = \bar{F}^{-1}(U)$. Transformations of $\bar{F}$, will generate appropriate biased samples. The first example of this, based on log-log transformation, is shown in Figure 6 (a).

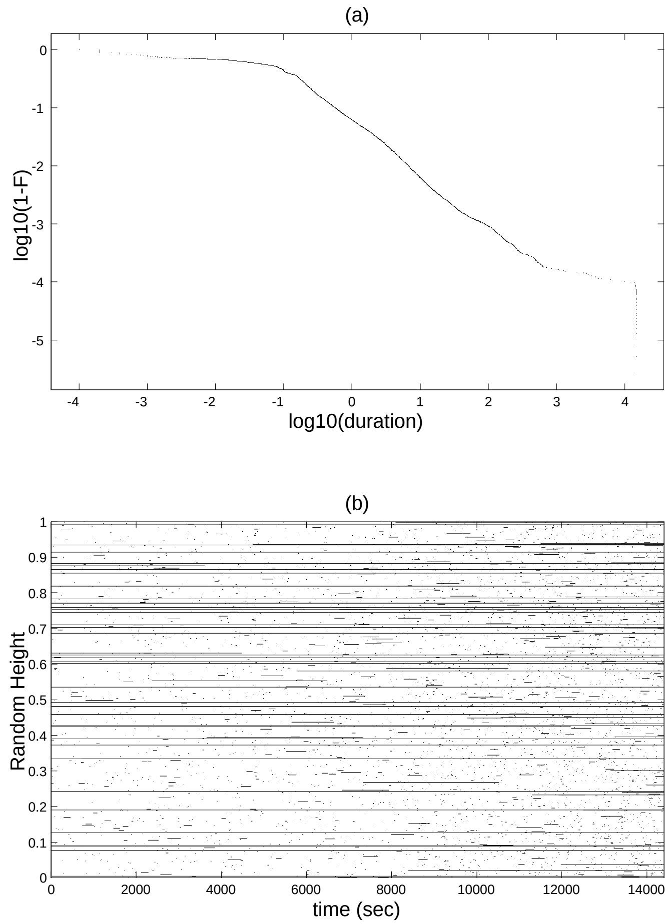

FIGURE 6: *(a) Log-log transformed CCDF, showing how Uniform sampling will put more weight on larger responses. (b) Corresponding Mice and Elephants plot, shows too much visual emphasis of the elephants.*

The horizontal axis in Figure 6 (a) shows the logarithm, with base 10, of the duration, and the vertical axis is the logarithm, with base 10, of the CCDF of the duration. This curve drives the sampling procedure as follows. First, select a random number $U \in [\log_{10}(1/N), 0]$, from the vertical axis and sample a corresponding response by choosing the maximum of the set

$$\{d_i : \log_{10} \bar{F}(d_i) \geq U\}, \tag{4}$$

where $d_i,\ i = 1, 2, \ldots, N$, represent the durations of the nonzero responses. Second, select another random number from the vertical axis and choose a corresponding response by (4). If the selected response is already in the sample, reject it and repeat the random sampling. Repeating these steps until we have 5000 responses completes the sample. Figure 6 (b) shows a Mice-Elephants plot of such a sample.

Figure 6 (b) shows too many elephants since the longer responses have a much higher probability of being chosen, because the log transformation puts very large weight there. This is the opposite of random sampling which obscured the elephants, because the untransformed CCDF put too little weight on them.

A family of compromises between these extremes is a sampling method based on the Box-Cox family of transformations (Box and Cox (1964)) of the CCDF of the durations. This family of transformations is defined as

$$G(d) = \frac{\bar{F}^{\lambda}(d) - 1}{\lambda \log 10}, \quad 0 \leq \lambda \leq 1. \tag{5}$$

The sampling procedure, after transforming the CCDF by (5), is the same as above. Note that, if $\lambda = 0$, then the sampling is the same as the one with a log_{10} transformation, and if $\lambda = 1$, then the sampling is consistent with random sampling. The results of applying one member of this family of transformations is shown in Figure 7.

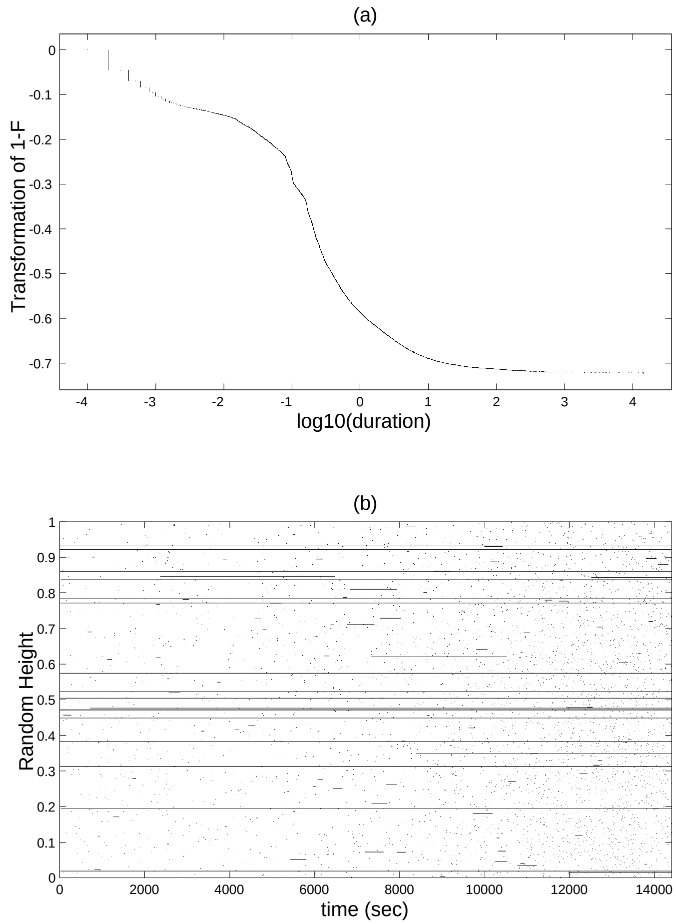

FIGURE 7: *(a) Box-Cox $\lambda = 0.6$ transformed CCDF, showing how Uniform sampling will distribute weights among response durations. (b) Corresponding Mice and Elephants plot, shows improved visual weighting.*

Figure 7 (a) shows the transformed CCDF of the duration with $\lambda = 0.6$. The horizontal axis is still the log_{10} of the duration. This shows that this Box-Cox transformation applies more weight to the intermediate size responses, than to the elephants, compared to the log_{10} transformation.

Figure 7 (b) shows the corresponding Mice-Elephants plot of the selected 5000 responses sampled using the transformed CCDF in Figure 7 (a). Several very long elephants appear, but they are fewer than those in Figure 7 (b). Also, the picture keeps many middle-sized responses.

The choice of $\lambda = 0.6$ was made on the basis of visual impression. Similar views, for a wide range of λ values are available in a movie, which is accessible in the Box-Cox Biased Sampling section of the web page González-Arévalo, et al (2004). The movie shows an increasing number of elephants in the plots as λ decreases from 1 to 0. The general choice of λ is an interesting open problem.

4 Quantile Window Sampling

Quantile window sampling provides a completely different view of the HTTP responses. The main idea is to display the HTTP responses, grouped according to their lengths of duration. This visualization enables us to focus on the structure of various subgroups of the mice and elephants that are similar in duration. We divide the population of responses, into consecutive subsets of size 5000, according to duration. That is, we display the longest 5000 responses, the next longest 5000 responses, and so on. It is interesting to study the succession of all such displays. A movie containing these can be found in the Quantile Window Sampling section of the web page González-Arévalo, et al (2004).

Figure 8 shows two such displays, at the elephant end of the scale.

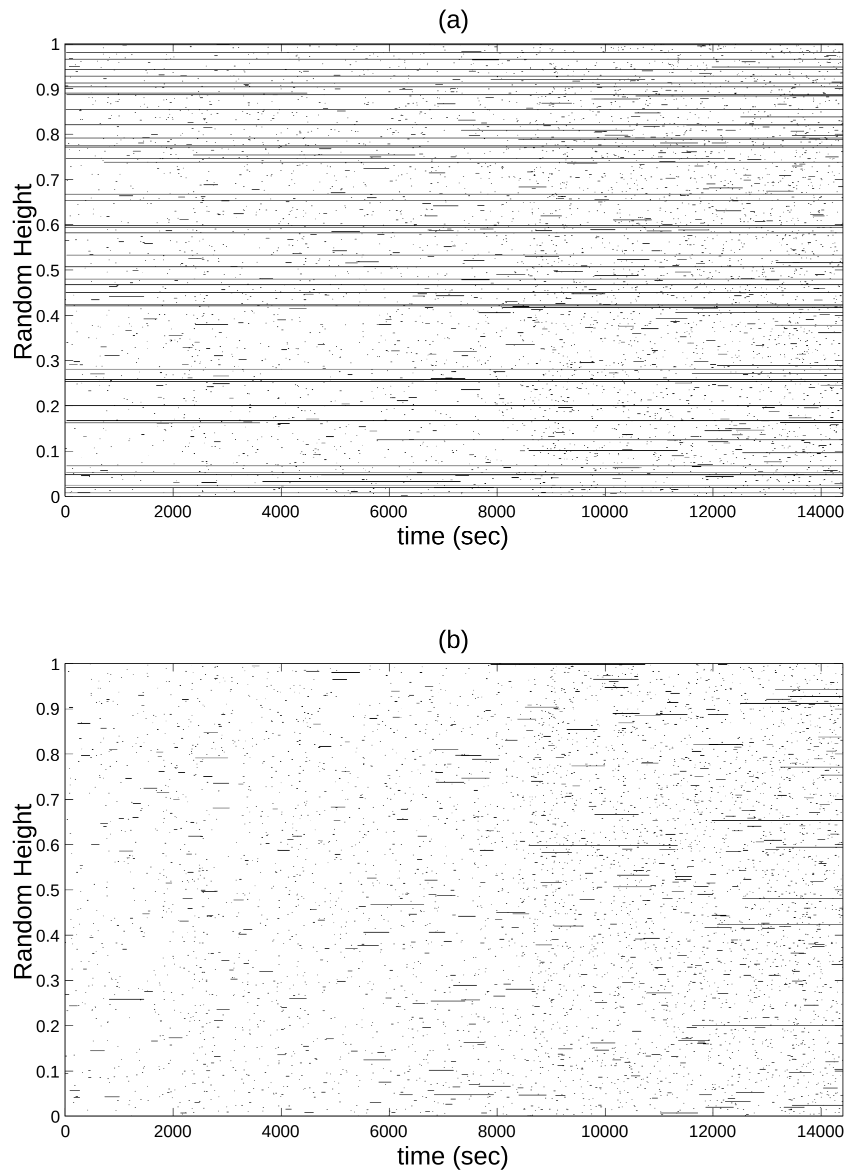

FIGURE 8: *For the same Sunday morning HTTP responses, (a) The 5000 largest elephants. (b) The biggest 50 Elephants are removed and the next biggest 5000 Elephants are displayed.*

Figure 8 (a) shows the biggest 5000 elephants of the population. As expected, the visual impression is dominated by the long elephants crossing over most of the full four hour time block. An interesting fact is that this group of responses contains a lot of mice also. To investigate this point further, we throw the biggest 50 elephants out and select the next biggest 5000 elephants from the population. Figure 8 (b) visualizes this sample. Only removing the biggest 50 elephants dramatically alters the visual impression, with now only middle-sized and short responses appearing. This allows a more clear view of other aspects of the population, including the fact that traffic level increases as the morning progresses, i.e. from left to right.

Figure 9 shows some lower quantile windows, chosen because they revealed some unexpected data artifacts.

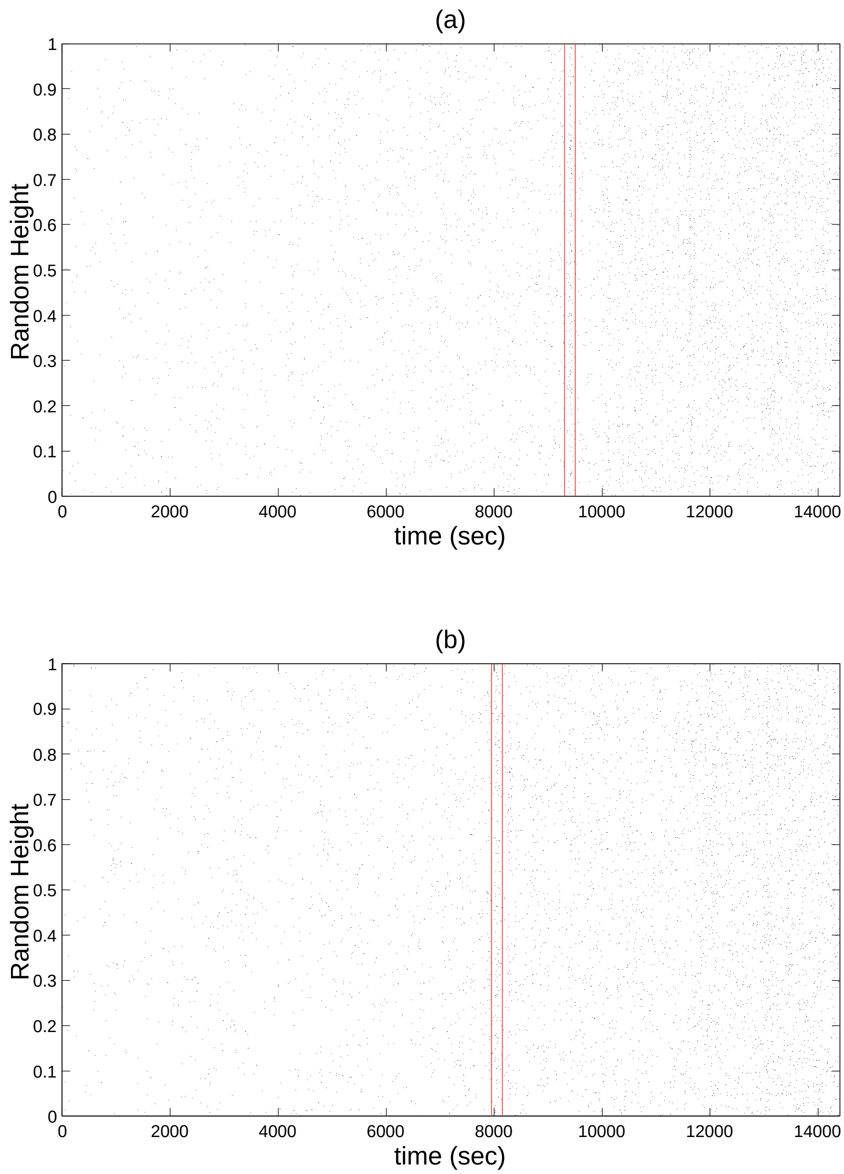

FIGURE 9: *Two plots displaying different quantiles of the population. Both show some vertical bursts in terms of starting times.*

Figure 9 (a) shows 5000 responses with duration from a minimum of 0.2146 to a maximum of 0.2262 (sec.). Basically, the responses shown in this figure have very similar duration lengths. An unexpected feature in this plot, are some high density vertical strips, representing bursts of responses of this duration at a few times. The biggest such burst appears around 9500 (sec.) and that region is marked by two vertical lines. More bursts in terms of the starting times are shown in Figure 9 (b). This plot shows 5000 responses with duration from a minimum of 0.6293 to a maximum of 0.7510 (sec.). The biggest such vertical strip is observed around 8000 (sec.) and that region is also marked by two vertical lines.

Insight into the causes of these bursts come from zooming in on the region between the red bars, and enhancing the visualization, as done in Figure 10. Additional useful information for understanding the causes of these bursts is the size of each HTTP response. Hence, in Figure 10, instead of using a random vertical height for separation, the size (in bytes) is used as the vertical axis.

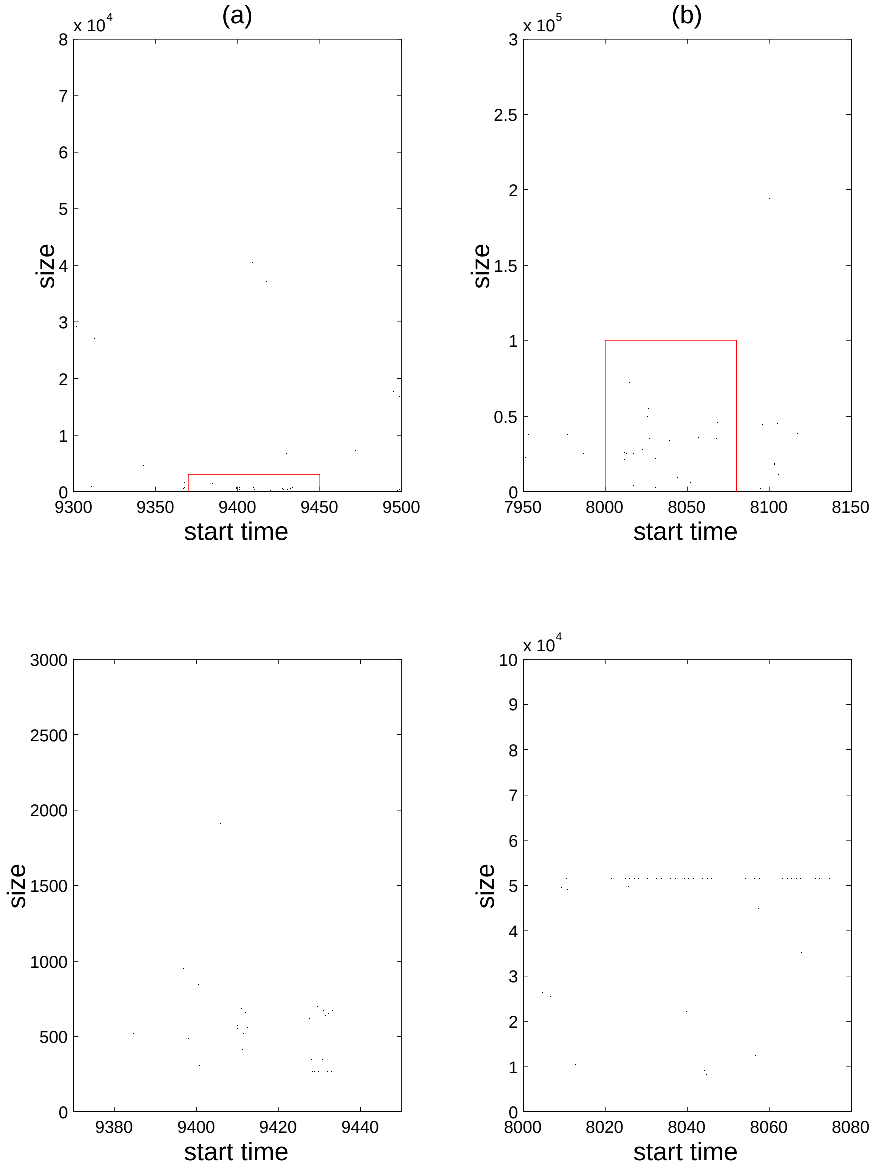

FIGURE 10: *The burst phenomena within two vertical lines in Figure 9 are deeply investigated by zooming in on the highlighted regions. The vertical axis in these plots is size of HTTP response, which aids the analysis. The left column shows zooms of Figure 9 (a), and the right column is zooms of Figure 9 (b). The bottom row is a further zoom, in on the rectangles drawn in the top row.*

The top panel of Figure 10 (a) contains the responses only within the two vertical lines in Figure 9 (a). Three big clusters are observed, with quite small response sizes and these are enclosed by a box. A further zoomed plot within this box is shown in the lower panel of Figure 10 (a). Around 9400, 9410, and 9430 (sec.), several small size responses are transferred almost at the same time, which creates the vertical strip that is seen in Figure 9 (a). A deeper investigation revealed that these responses were composed of two packets, and came from a common server, suggesting the download of a number of small web page components, such as embedded thumbnail images. The three clusters have a timing that suggests they represent three such web page. Usually such transfers involve single packets, and thus are not highlighted in the present analysis, because single packet responses have 0 duration. The reason that these responses come as two packets appears to be that the server sent the header of the response object in one packet, and then, after some extra processing, sent the actual data. HTTP attaches an application-level header to each response, and it seems that the particular server sending those responses is splitting header and data, and also forcing TCP to send the header as soon as it is ready. This uncommon behavior has some performance benefits. The main benefit is that the client can start processing the header right away, rather than waiting for the header plus the object to arrive, slightly reducing the user perceived delay.

The top panel of Figure 10 (b) shows another type of burst. This plot contains the responses only within the two vertical lines in Figure 9 (b). This burst of HTTP responses is far different from the burst of embedded small objects discovered in Figure 10 (a). This time it is seen that the burst of responses all have essentially the same size. This strong similarity of size is confirmed by zooming in on the rectangular box, as shown in the lower panel. These responses are large enough that were sent as multiple packets. Deeper investigation showed that these all came from the same server, which was an Internet game site. These appear to be a sequence of updates of game states, which explains the common size.

5 Commonality of flow rates

Another interesting view, of the HTTP response data analyzed in Section 2, is a scatterplot of the duration (time, i.e. length of the line segments) of each response, versus the size of the response in bytes (i.e. the amount of data transferred). Both variables share the heavy tailed *mice and elephants* behavior demonstrated in Figures 1-4, so a reasonable view of the data comes from plotting both variables on the log scale. Figure 11 shows the resulting scatterplot. This requires special handling of responses with 0 duration (e.g. this happens for single packet responses). This was done by dropping such responses from the sample, which resulted in 382,127 responses appearing in the scatterplot.

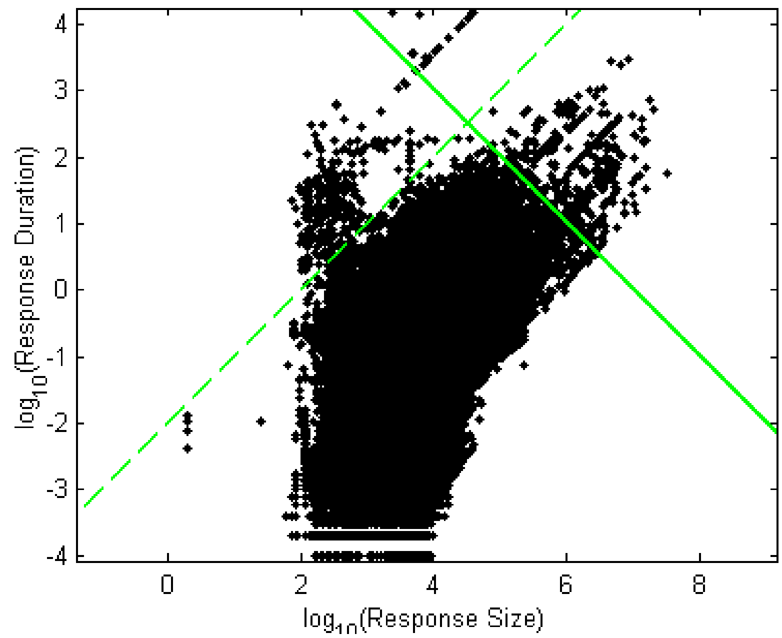

FIGURE 11: *Log - log scatterplot showing how flow duration time depends on transmission size. Suggests clusters of flows with same throughput (rates), as diagonal lines.*

The general tendency in Figure 11 is roughly what one might expect: larger size responses need more time, so there is a general upwards trend. Horizontal lines at the bottom of the plot reflect discreteness of very small time measurements. A perhaps surprising feature is the diagonal lines of points present at larger times and sizes. Not only do the lines appear to be parallel, they also lie at a 45° angle to the coordinate axes, as indicated by the parallel dashed green line, which has equation $y = x - 2$. These diagonal lines of points represent sets of flows where

$$\log_{10} time = \log_{10} size + C,$$

for some constant C, which is the same as

$$size = R \cdot time,$$

where $R = 10^{-C}$ is interpretable as a *constant rate*. Thus the flows following each diagonal line have essentially the same rate (defined as total size divided by total time).

Figure 11 gives a strong visual impression that the large flows may be *naturally clustered* in terms of rates. This is sensible because rates are naturally driven by the nature of the network between the source and the destination. Most of the computers within UNC will likely have quite similar rates to a few popular web-sites, resulting in similar rates for large numbers of transfers.

The second main open problem of this paper is to develop methods for analyzing this aspect of the population. How can the clusters be isolated? What are the cluster boundaries? How many flows are in the major clusters?

A start on addressing these issues appears in Figure 12. Here the data are projected onto the orthogonal solid green line in Figure 11, so the problem is reduced to studying clusters in univariate data. For easy visual connection to Figure 11, the data are transformed to the coordinate system which allows treating the solid green line as the axis. In particular the transformation is:

$$proj = -2 - \frac{\log_{10} time - \log_{10} size}{\sqrt{2}}.$$

The denominator of $\sqrt{2}$ makes the transformation length invariant (i.e. a rotation), and the subtraction from -2 gives the most straightforward view of the solid green line as an axis. In particular, the solid green line is rotated in a counterclockwise fashion to give a conventional horizontal axis (as in Figure 12), i.e. the points on the upper left part of the green line, become negative values in the univariate view.

The top panel of Figure 12 shows two displays of the projected data. The first is the green dots, which are a standard *jitter plot* (again see Tukey and Tukey (1990) and Cleveland (1993)), where the horizontal coordinate is the projection value (i.e. location of each data point when projected onto the green line), and a random vertical coordinate is used for visual separation. The jitter plot shows only a random sample of 10,000 to avoid overplotting problems. The second display of the data is the family of blue curves. These are kernel density estimates (essentially smooth histograms), with a wide range of window widths. Looking at a family of smooths is the *scale space view* of data, which is recommended as a practical solution to the traditional problem of bandwidth choice, see Chaudhuri and Marron (1999) for further discussion.

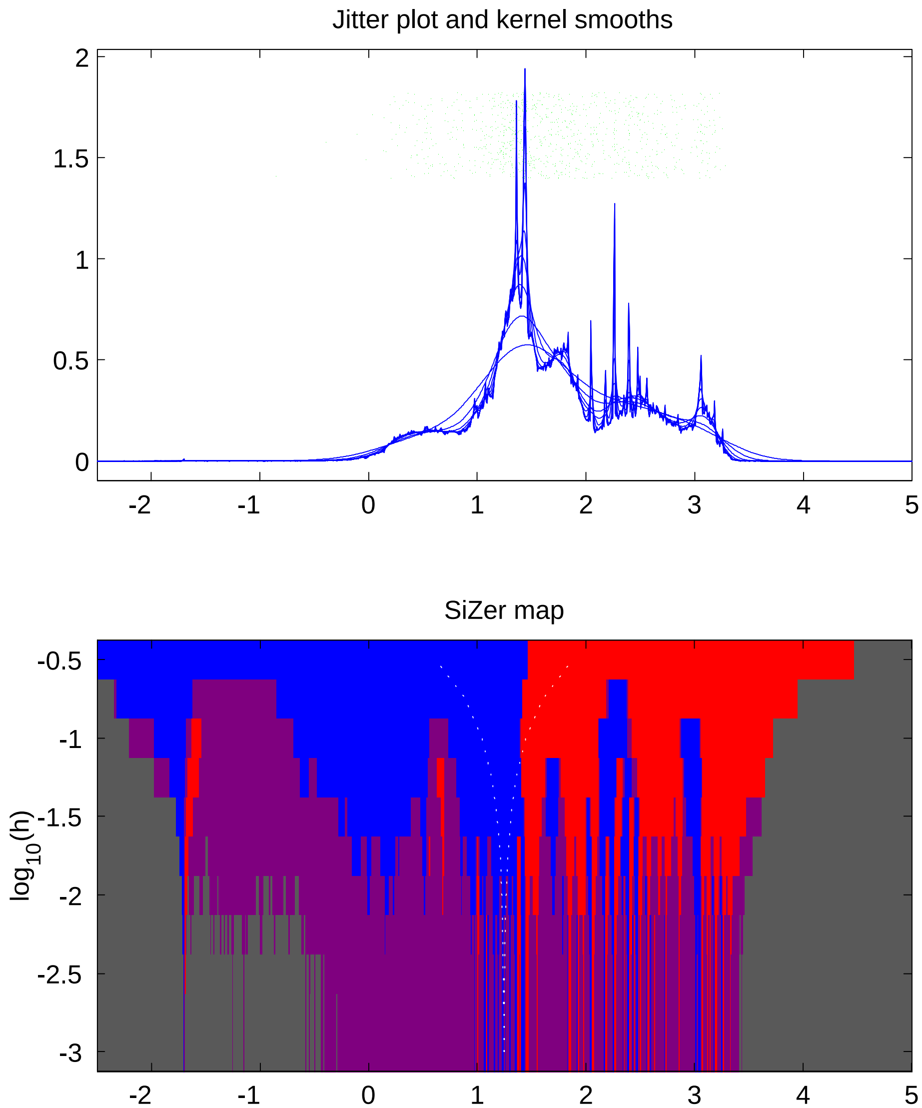

FIGURE 12: *SiZer analysis of projected scatterplot. Shows many significant clusters.*

The family of kernel smooths suggests a number of *broad bumps*, and there are also a number of *small spikes*. It is tempting to dismiss the spikes as *spurious sampling variability*, however recall that such clusters were suggested in Figure 11, and a possible physical explanation was suggested above. Furthermore the sample size $N = 382,127$ is fairly large, so perhaps those spikes represent important underlying structure in the data?

A useful tool for addressing such exploratory data analysis questions is the SiZer map shown in the bottom panel of Figure 12. Rows of this map correspond to different window widths, i.e. to blue kernel smooths, and the horizontal axis is the same as in the top panel. Colors are used to indicate statistical significance of the slopes of the blue curves, with blue (red, resp.) for significantly increasing (decreasing, resp.), with purple for regions where the slope is not significantly different from 0, and with gray where the data are too sparse for reliable inference.

The SiZer map shows that all of the *broad bumps* are statistically significant, as are most of the tall thin bumps. These may not be surprising because $N = 382,127$ allows resolution of quite a few features of the underlying probability density, in view of the large sample size. More surprising may be the very small bump at -1.7. This is hardly visible in the blue family, and yet is clearly statistically significant in the SiZer map.

The analysis of Figure 12 is not very satisfying, because it seems that perhaps some of the clusters, that are clearly visible as lines of points in Figure 11, might be *masked* by the large amount of other data that makes up the broad peaks. A simple approach to this is to repeat the analysis for a suitably thresholded sub-sample. Visual inspection of Figure 11 suggests using only the data above the solid green line, $y = -x + 7$. There were 572 such points, still enough for effective kernel density estimation. The resulting analysis is shown in Figure 13.

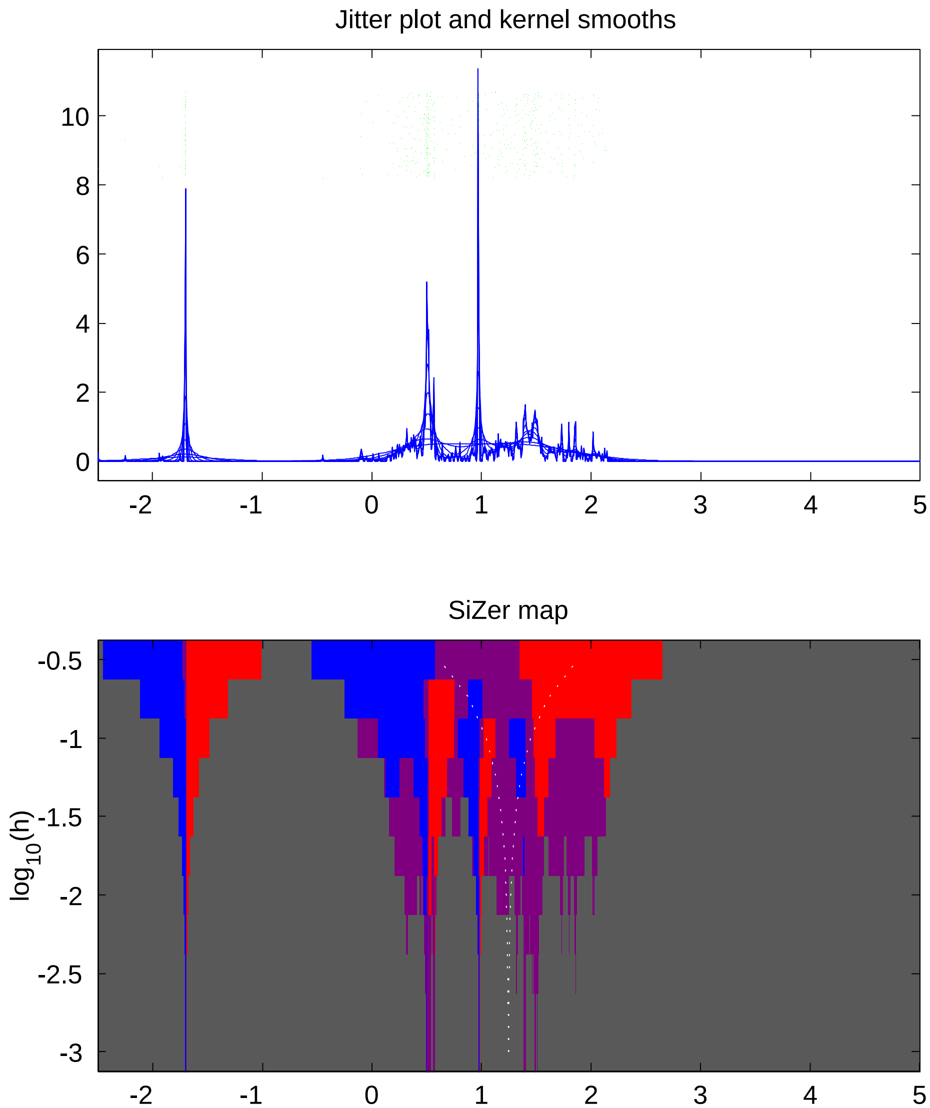

Figure 13: *SiZer analysis of scatterplot points above the solid green line in Figure 11. Shows significant clusters, but different from those of Figure 12. Thus many more clusters than those in Figure 11 exist.*

As expected, the *broad bumps* in Figure 12 have now disappeared. There are also some very significant *slim spikes*. Note that the spike near -1.7 is now much taller in the blue family of curves (essentially all of these data points have been retained from Figure 12, and now are proportionally a far larger part of the population). However, note that many of the tall thin peaks in Figure 12 are not present in Figure 13. This shows that much of the *clustered aspect* of the population actually occurs more in the main body of the main scatterplot in Figure 11, and thus can not be teased out by simple thresholding as done in Figure 13.

Thus, this is a case where the scatterplot of Figure 11 hides a large amount of interesting population structure. The SiZer analysis is an indirect way of understanding this. Are there more direct ways of visualizing this type of structure?

6 Acknowledgements

This collaboration was a result of the authors meeting in the SAMSI program on Network Modeling for the Internet. The third author is grateful to the Cornell University School of Operations Research and Industrial Engineering for hospitality, support, and an exciting research environment. This research was partially supported by Cornell University's College of Engineering Mary Upson Fund and by NSF Grants DMS-9971649 and DMS-0308331.

References

[1] Box, G. P. E. and Cox, D. R. (1964). An analysis of transformations. *Journal of the Royal Statistical Society, Series B* 26, 211–243.

[2] Cao, J., Cleveland, W. S., Lin, D. , and Sun, D. X. (2001) On the Nonstationarity of Internet Traffic, *Proceedings of the ACM SIGMETRICS '01*, 102-112. Internet available at: http://cm.bell-labs.com/cm/ms/departments/sia/wsc/webpapers.html.

[3] Cao, J., Cleveland, W. S., Lin, D. , and Sun, D. X. (2002a) The Effect of Statistical Multiplexing on the Long Range Dependence of Internet Packet Traffic, Bell Labs Tech Report, 2002, Internet available at: http://cm.bell-labs.com/cm/ms/departments/sia/wsc/webpapers.html.

[4] Cao, J., Cleveland, W. S., Lin, D. , and Sun, D. X. (2002b) Internet Traffic: Statistical Multiplexing Gains, *DIMACS Workshop on Internet and WWW Measurement, Mapping and Modeling*, Internet available at: http://cm.bell-labs.com/cm/ms/departments/sia/wsc/webpapers.html.

[5] Cao, J., Cleveland, W. S., Lin, D. , and Sun, D. X. (2002c) Internet Traffic Tends Toward Poisson and Independent as the Load Increases, *Nonlinear Estimation and Classification*, eds. C. Holmes, D. Denison, M. Hansen, B. Yu, and B. Mallick, Springer, New York, Internet available at: http://cm.bell-labs.com/cm/ms/departments/sia/wsc/webpapers.html.

[6] Chaudhuri, P. and Marron, J. S. (1999) SiZer for exploration of structure in curves, *Journal of the American Statistical Association*, 94, 807-823.

[7] Cho, K. (2001) Aguri Traffic Density Graph, http://www.csl.sony.co.jp/person/kjc/kjc/software/aguri-density.html

[8] Cleveland, W. S. (1993) *Visualizing Data*, Hobart Press, Summit, New Jersey, U.S.A.

[9] Cox, D. R. (1984) Long-Range Dependence: A Review, in *Statistics: An Appraisal, Proceedings 50th Anniversary Conference*. H. A. David, H. T. David (eds.). The Iowa State University Press, 55-74.

[10] Crovella, M. E. and A. Bestavros, A. (1996) Self-similarity in world wide web traffic evidence and possible causes, *Proceedings of the ACM SIGMETRICS 96*, pages 160–169, Philadelphia, PA.

[11] Downey, A. B. (2000) The structural cause of file size distributions, Wellesley College Tech. Report CSD-TR25-2000. Internet available at: http://rocky.wellesley.edu/downey/filesize/.

[12] Downey, A. B. (2001) Evidence for long tailed distributions in the internet, ACM SIGCOMM Internet Measurement Workshop, November 2001. Internet available at http://rocky.wellesley.edu/downey/longtail/.

[13] Duffield, N. G. and Lund, C. (2003) Predicting Resource Usage and Estimation Accuracy in an IP Flow Measurement Collection Infrastructure, *ACM SIGCOMM Internet Measurement Conference 2003*.

[14] Gong, W., Liu, Y., Misra, V. and Towsley, D. (2001) On the tails of web file size distributions, *Proceedings of 39-th Allerton Conference on Communication, Control, and Computing*. Oct. 2001. Internet available at: http://www-net.cs.umass.edu/networks/publications.html.

[15] González-Arévalo, B., Hernández-Campos, F., Marron, J. S. and Park, C. (2004) Internet site: http://www.cs.unc.edu/Research/dirt/proj/marron/VisualChallenge/

[16] Hannig, J., Marron, J. S. and Riedi, R. (2001) Zooming statistics: Inference across scales, *Journal of the Korean Statistical Society*, 30, 327-345.

[17] Heath, D., Resnick, S. and Samorodnitsky , G. (1998) Heavy tails and long range dependence in on/off processes and associated fluid models, *Mathematics of Operations Research*, 23, 145-165.

[18] Hernández-Campos, F., Marron, J. S., Samorodnitsky, G. and Smith, F. D. (2002) Variable Heavy Tailed Durations in Internet Traffic, to appear in *Performance Evaluation*, web available at http://www.cs.unc.edu/Research/dirt/proj/marron/VarHeavyTails/.

[19] Hohn, N. and Veitch, D. (2003) Inverting Sampled Traffic, *ACM/SIGCOMM Internet Measurement Conference*, 222-233.

[20] Karagiannis, M. F. T., Molle, M. and Broido, A. (2004) A Nonstationary Poisson View of Internet Traffic, in *Proceedings of IEEE Infocom*.

[21] Leland, W. E., Taqqu, M. S., Willinger, W. and Wilson, D. V. (1994). On the Self-Similar Nature of Ethernet Traffic (Extended Version), *IEEE/ACM Trans. on Networking*, 2, 1-15.

[22] Mandelbrot, B. B. (1969) Long-run linearity, locally Gaussian processes, H-spectra and infinite variance, *International Economic Review*, 10, 82-113.

[23] Marron, J. S., Hernández-Campos, F. and Smith F. D. (2002) Mice and Elephants Visualization of Internet Traffic, *Proceedings of the CompStat 2002*, Härdle, W., Rönz, B. (Eds.), Springer, Berlin, 47-54, internet available at: http://www.cs.unc.edu/Research/dirt/proj/marron/MiceElephants/.

[24] Reed, W. J. and Jorgensen, M. (2004) The double Pareto-lognormal distribution - A new parametric model for size distribution. *Communications in Statistics -Theory & Methods*, 33, 1733-1753.

[25] Resnick, S. and Samorodnitsky, G. (1999) Activity periods of an infinite server queue and performance of certain heavy tailed fluid queues, *Queueing Systems*, 33, 43-71.

[26] Smith, F. D., Hernández, F., Jeffay, K. and Ott, D. (2001) "What TCP/IP Protocol Headers Can Tell Us About the Web", *Proceedings of ACM SIGMETRICS 2001/Performance 2001*, Cambridge MA, June 2001, pp. 245-256.

[27] Taqqu, M. and Levy, J. (1986) Using renewal processes to generate LRD and high variability, in: *Progress in probability and statistics*, E. Eberlein and M. Taqqu eds. Birkhaeuser, Boston, 73-89.

[28] Tukey, J., and Tukey, P. (1990). Strips Displaying Empirical Distributions: Textured Dot Strips. Bellcore Technical Memorandum.

[29] Zhang, Z.-L. Ribeiro, V. J., Moon, S. B. and Diot, C. (2003) Smalltime scaling behaviors of Internet backbone traffic: an empirical study, in *Proceedings of IEEE Infocom.*