

Capturing the Variability of Internet Flows Across Time

Jörg Wallerich

Anja Feldmann

Technische Universität München, Germany

Abstract—More and more traffic management techniques, including accounting and load adaptive routing, try to take advantage of the fact that traffic demands are consistent with Zipf’s law. By treating a few large volume demands differently they try to capture most of the traffic. This relies on the implicit assumption that traffic demands are persistent in volume over time; meaning that their volume does not change drastically over time.

As this assumption has been shown to be incorrect we in this paper focus on how Internet flows behave over time. Accordingly, this paper examines the characteristics of volatility in a qualitative way by characterizing the components that are responsible for changes in the cast of heavy hitters over time.

I. INTRODUCTION

Starting with the statistical analysis of Ethernet LAN traces in [1], there has been extensive work towards developing appropriate mathematical and statistical techniques for explaining, describing, and validating the invariant that data traffic is consistent with self-similar behavior, e.g., [2], [3], [4], [5], [6], [7]. This implies that there is significant variability in the overall traffic volume. Another area of extensive work concerns the analysis of the frequency of occurrences of an event. If the frequency of occurrences as function of the rank is consistent with a power-law distribution it is referred to as Zipf’s-like (see [8], [9] and references therein). The rank is determined by the frequency of the occurrence of the studied event, where a low rank index refers to a popular event. Not surprisingly quite a number of different quantities in Internet traffic are consistent with Zipf-like distributions, including the popularity of Web pages [10], [11], traffic demands [12], [13], [14], [15], [16], as well as interdomain Web traffic demands [17]. This raises the question if it is possible to take advantage of this invariant. If data is consistent with a Zipf-like distribution, it is in principle sufficient to concentrate on the popular events when engineering the network in order to optimize the overall performance. This is the common idea, i.e., used for measurement [18], traffic engineering [12], [19], [13], [20], including load adaptive routing [21], scheduling [22], Web caching [10], [11], [23], etc..

All of the above examples assume and rely on the assumption that if an event is popular it will stay popular. On the other hand it is well-known that traffic is variable and traffic fluctuations can induce significant changes in popularity. Whether this has an impact on the above applications depends on two rates: the *rate of event popularity changes* and the *engineering rate*, the rate at which engineering reacts to the changes in

popularity. In this paper we examine the characteristics of the instabilities in the context of traffic flows, building on the work of Papagiannaki et al. [24] and Wallerich et al. [15].

Traffic flows are the result of grouping packets using different strategies ranging from the most specific scheme of using the full five-tuple of source and destination IP-addresses and port numbers together with the protocol, through schemes that use BGP based destination prefixes, to schemes that map the IP addresses to ASes and use these as destinations, e.g., [25], [26].

In general instabilities can have many causes including bandwidth fluctuations. But even if one assumes that all flows operate at a fixed bandwidth, a top ranked flow can become a medium ranked one if a number of larger flows pop up. Alternatively a medium ranked flow can become a highly ranked one if top ranked flows end or reduce their bandwidth. In this paper we examine their characteristics and implications.

The remainder of the paper is organized as follows: in Section II we give an overview over our datasets and explain how we prepare the packet level and NetFlow [27] traces to derive the per flow statistics and rankings for our study. Section III presents our results regarding the components of the instabilities. We summarize our experience in Section IV.

II. TRACES AND DATA PREPARATION

A. Traces

In this paper we use two traces. The first trace (MWN), a 24 hour packet level trace, was collected on Nov. 15., 2005 in Munich, Germany at the *Munich Scientific Research Network* (Münchner Wissenschaftsnetz, MWN) which connects two major universities, some colleges, and affiliated research institutes to the Internet, totaling to approximately 50,000 hosts. The volume transferred over its Gigabit Internet link is around 2 TB a day and the average utilization during busy-hours is about 350 Mbps (68 K packets per second).

The second trace (WASHng), a 24 hour NetFlow trace, was collected with a sample rate of 1/100 on Nov. 2., 2005 on a core router in Washington, USA, of the Abilene [28] network.

B. Flow Generation and Aggregation

The packet level trace MWN is used to generate standard 5-tuple NetFlow records. NetFlows are terminated if they are idle (receive no packet) for longer than an inactivity timeout of 15 seconds (default value for Cisco’s NetFlow [27]). For each flow, we derive start and end times as well as byte and

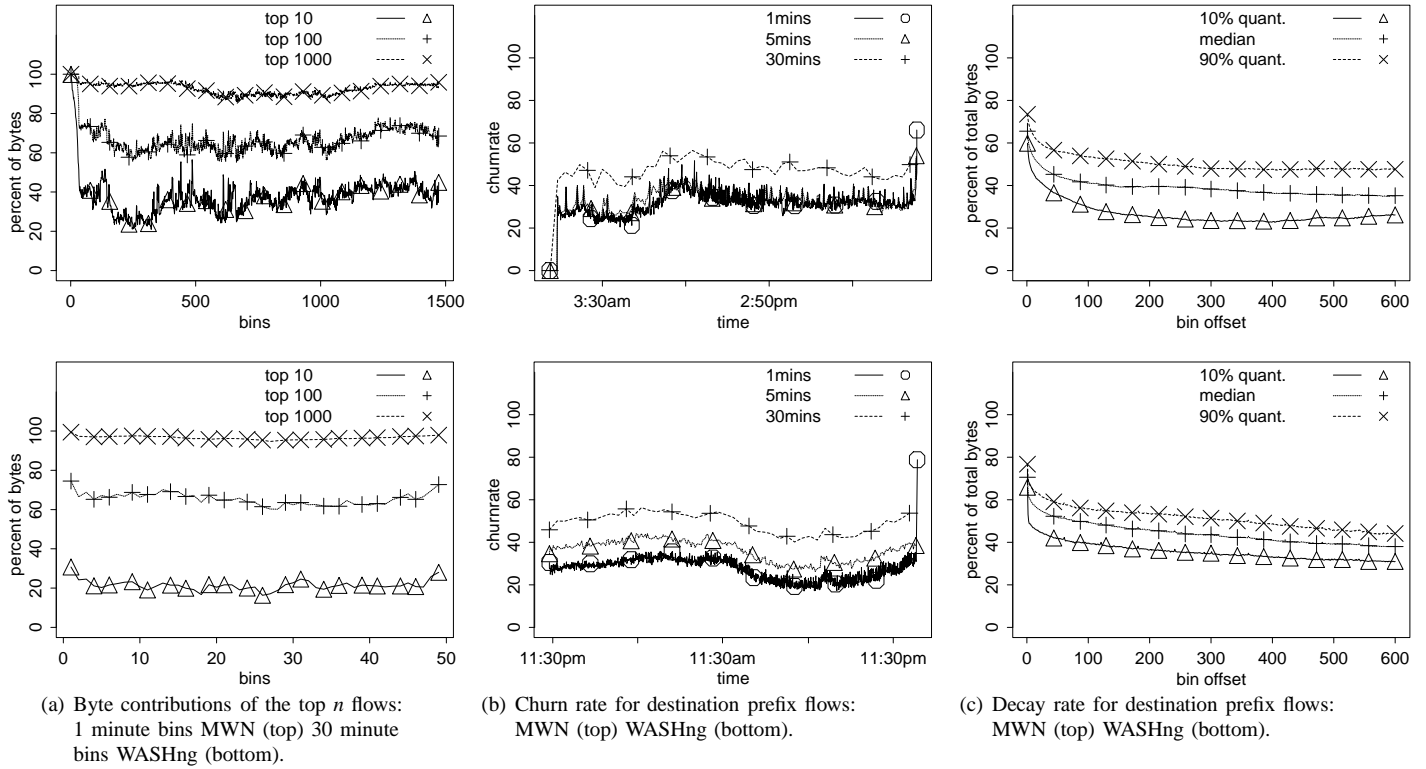


Fig. 1. Basic Data Characteristics

packet counts. As traffic engineering is usually not done on 5-tuple flows, we generated aggregate flows based on destination prefix from the 5-tuple flows. Prefix aggregation is done by accumulating bytes and packets for flows that have the same destination prefixes, while using the same criteria for flow termination as for non-aggregated flows. For the MWN packet level trace a BGP table extracted from the MWN access router was used to determine the prefix length masks. For the WASHng NetFlow trace, the exporting router includes its prefix length masks in the NetFlow records. While 5-tuple flows have constant average bandwidth information, the bandwidth of destination prefix flows depends on the number and the bandwidth of the 5-tuple flows that contribute to an aggregated flow at any given time. We compute the bandwidth per destination flow over 1, 5, and 30 minute time periods.

Since the MWN packet trace is collected on an access router the address variability on the access (MWN) side is limited. Therefore we separate the MWN trace in outgoing and incoming traffic and only consider destination/source prefix aggregation for outgoing/incoming traffic. From the 3.3 billion packets we create 9.1 million destination prefix flows for MWN and from the 11.2 million NetFlows 1.9 million destination prefix flows for WASHng.

C. Ranking and Zipf's law

We next compute flow rankings for each time period, also referred to as time bins, based on the destination prefix flows covering that bin. The ranking is determined based on the

contribution in terms of bytes of each flow to the time bin. A ranking is an ordered set R of flows f with

$$r_i < r_j \Leftrightarrow \text{bytes}(r_i) > \text{bytes}(r_j)$$

where r_i is the rank of flow f_i and $\text{bytes}(r_i)$ is the bytes of the flow fragment for time bin i . $r_i = 1$ represents the flow that contributes the most bytes and therefore has the highest rank. In terms of bin sizes we used time granularities of 1, 5 and 30 minutes to compute rankings of aggregated flows.

Zipf's law for flow rates states that for a set of n inferred flow rates, ordered as $x_{(1)} \geq x_{(2)} \geq \dots \geq x_{(n)}$, we may think of $x_{(r)}$ as the r th-largest flow rate and of r as the flow rate's rank ($1 \leq r \leq n$), the relationship $rx_{(r)} = \text{constant}$ (or, more generally, $r^\alpha x_{(r)} = \text{constant} = c, \alpha > 0$) holds, at least approximately. If flow rate data is consistent with Zipf's law it implies that the top-ranked flow rates are exceptionally large but rare and the lower-ranked rates are smaller but more common. This follows directly from the *size-frequency* relationship that corresponds to Zipf's law and states that $f(x)$, the relative frequency of occurrence of a flow rate of size x satisfies the relation $f(x) = c \cdot x^{-2}$ (or, more generally, $f(x) = c \cdot x^{-(1+\alpha)}, x = 1, 2, \dots$).

III. CHARACTERIZATION AND ANALYSIS

In this section we explore possible causes for instabilities in the rankings of large flows. After all Zipf's law in itself leaves open the possibility that the cast of flows that contribute can

vary considerably over time. For this purpose we explore the characteristics of the processes contributing to the instabilities: arrival of new flows, departure of old flows, and bandwidth changes.

A. Basic data characteristics

We start by verifying that our datasets are indeed consistent with Zipf’s law even for different bin sizes and across time by exploring the popularity of destination prefixes for multiple consecutive bins. For this purpose we rank destination prefix flows by the total traffic they sent out of the MWN during a certain time period from largest to smallest, and plot the percentage of total bytes from largest to smallest. As expected we find (Figure not shown) a linear relationship on the log-log scale, an indication that the distribution is consistent with the characteristics of a Zipf-like distribution. This is the case for arbitrary time periods, different time granularities, and also for the WASHng data set confirming previous results [24], [15]

While these plots indicate that the data is consistent with Zipf’s law they also show that the volume of the top flows changes over time. To explore the overall impact of the bandwidth changes of flows we, in Figure 1(a), examine what percentage of the outgoing traffic is accounted for by the top 10, top 100, and top 1000 destination prefix flows. Note, that the top 100 flows already make up for more than 60% of all bytes, independent of the type of trace and the bin length. A typical 1, 5, 30 minute bin for destination prefixes contains about 106,000, 485,000, 2.8 million flows on average for MWN and 33,000, 152,000, 897,000 for WASHng.

Figure 1(a) leaves open that the cast of popular flows can vary considerably from one time period to the next. Indeed if one considers aggregated flows one can expect their bandwidth utilization to be consistent with self-similarity behavior [7]. The impact of the bandwidth variability is highlighted by Figure 1(b) which shows the “churn rate” among the top 1000 flows. Here the churn rate is defined to be the percentage of top n flows in that bin that were not among the top n flows in the previous bin. Figure 1(b) shows, that indeed a significant portion of top 1000 ranked flows are no longer in the top 1000 flows in the next bin. About 30% of the cast of top 1000 flows is no longer top 1000 in the next bin. One reason for this effect is due to the arrival of new flows, that make it into the top 1000 ranks and the termination of old top 1000 flows. One interesting aspect here is that the churn rate increases with larger time bins, as it gets more and more difficult for a flow to retain a top 1000 rank over longer time periods. This change in the top ranks results in the undesired fact, that a set of flows once classified as top n is at first responsible for a significant portion of the overall traffic, but that this portion quickly diminishes over time. This confirms previous observations [24], [15] for different data sets.

In Figure 1(c) we plot the 10%, 50% and 90% quantiles of the fraction of bytes contributed by a fixed set of top 100 flows over time. The plot shows, that for any such set, the median of the byte contributions lies at first between 60% and 75%, but

within about 20 minutes, this fraction drops to about 40%. For 10% of the flows it even drops below 30%. These observations are consistent with findings by Papagiannaki et al. [24] where the authors show a similar decrease in byte coverage.

B. Flow Entry and Exit Process

The observation that the fraction of bytes captured by the top n flows drops quickly motivates us to explore the reasons for such behavior. Accordingly we go beyond the previous analysis [24], [15] and decompose the churn rate into its components, flow entry process, flow exit process, and duration, and study their behavior. The entry process determines the rank at which a flow enters the ranking and the flow departure process determines the rank at which it leaves the ranking. The exit process is the result of combining the entry process with the durations. Due to space constraints we, in this paper, only focus on the flow entry and exit processes. After all if a flows enters (leaves) the top n flows, then one of the old flows has to be moved out from (moved into) the top n flows.

We start with the entry process. Yet, instead of immediately jumping into the data analysis we first explore what to expect. Accordingly, we ask the question, with what probability will a new flow pop up in the top n flows. This corresponds to the question if a newly arriving flow contributes more bytes than at least one of the old top n flows, that is its byte contribution is in the interval $]bytes(r_{i+1});bytes(r_i)[$. This is equivalent to the question, with what probability falls a sample from a random variable (in this case, a flow bandwidth) into a certain target interval. For a random variable \mathbf{X} of a given probability distribution function P , we have $\mathbf{X} = P^{[-1]}(\mathbf{U})$, where \mathbf{U} is a random variable that is uniformly distributed in $[0; 1]$. This relationship is commonly used to generate random values of a specific probability distribution. We can now estimate the probability of a new sample s to have a value between any two elements s_i and s_{i+1} of S by determining the probability that $P(s)$ falls between the corresponding two elements of \bar{S} . In effect P works as a projection from S into $[0; 1]$ and thus allows us to use the properties of the uniform distribution to conveniently determine the probability for a new sample to fall in a given interval. As the elements in \bar{S} are uniformly distributed in $[0; 1]$ any new element chosen uniformly at random in $[0; 1]$ is likely to fall between any two consecutive elements of $\langle 0, \bar{S}, 1 \rangle$ with the same probability and therefore between any two consecutive elements of $\langle P^{[-1]}(0), S, P^{[-1]}(1) \rangle$. This implies that one should expect that the distribution of the entry ranks is consistent with a uniform distribution independent of the distribution from which the original samples are drawn.

In Figure 2 (left and middle columns) we show histograms of the first rank of those flows that have an average rank of at least 100. Contrary to the theoretical considerations, these do not appear to be consistent with a uniform distribution. Still, the plots indicate that new flows arrive with all kinds of different ranks including the top ones even though there is a

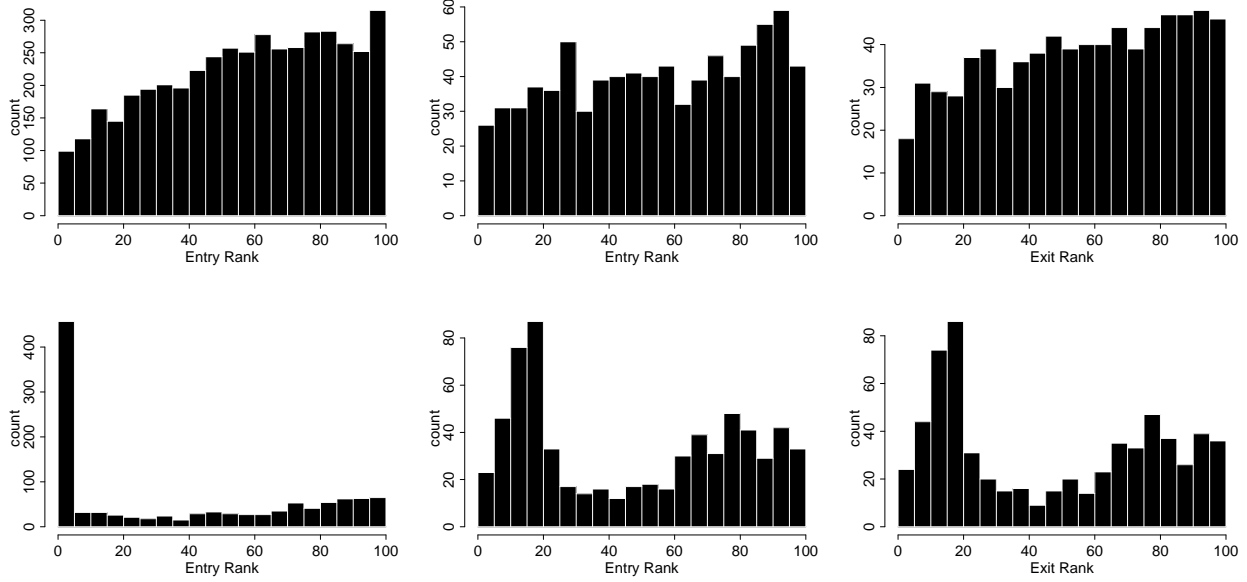


Fig. 2. Histogram of entry/exit ranks: top MWN, bottom WASHng; left entry 1 minute bins, middle entry 30 minute bins, right exit 30 minute bins.




Fig. 4. Density of the logarithm of the flow duration for WASHng.

bias towards larger ranks for the MWN trace. For the WASHng trace there is a bias towards top 5 ranks for 1 minute bins and the top 20 ranks for 30 minutes bins. Most likely this bias is due to the kind of traffic on the Abilene network which is mainly due to research environments. For both data sets the distribution of data becomes more even across the bins (closer to a uniform distribution) as we consider larger bin sizes.

Recalling the theoretical analysis of the entry process, we note that the properties of the process, that removes one of the top n samples instead of adding it, are the same. Therefore the distributions should be the same. Figure 2 (right column) shows a histogram over the ranks of top 100 flows, when they terminate. Again the data is not consistent with a uniform distribution. Yet, the exit processes show almost the same distribution as the entry processes and move closer to a uniform distribution as one increases the bin sizes or excludes the top most ranks. This holds for 1, 5 and 30 minute bins and both traces.

The above results for the flow entry and exit processes show that newly arriving flows and departing flows are with a non-negligible probability top ranked. This causes a lack of predictability and therefore it is not obvious how to take advantage of the “heavy-tailed” nature of flow rates for the purpose of better engineering the network. Still, Papagiannaki et al. [24] found that they were able to identify a subset of flows that contribute more traffic than the average flows, i.e., they are highly ranked, and persist longer than the typical flows. For this to be possible some of the highly ranked flows have to have longer durations. We find that using top rank of a flow (its highest rank) is not a good indicator. But flows with top average ranks (over their duration) usually persist longer than those with lower average ranks, see Figure III-B. These findings explain the observations by Papagiannaki et al. [24].

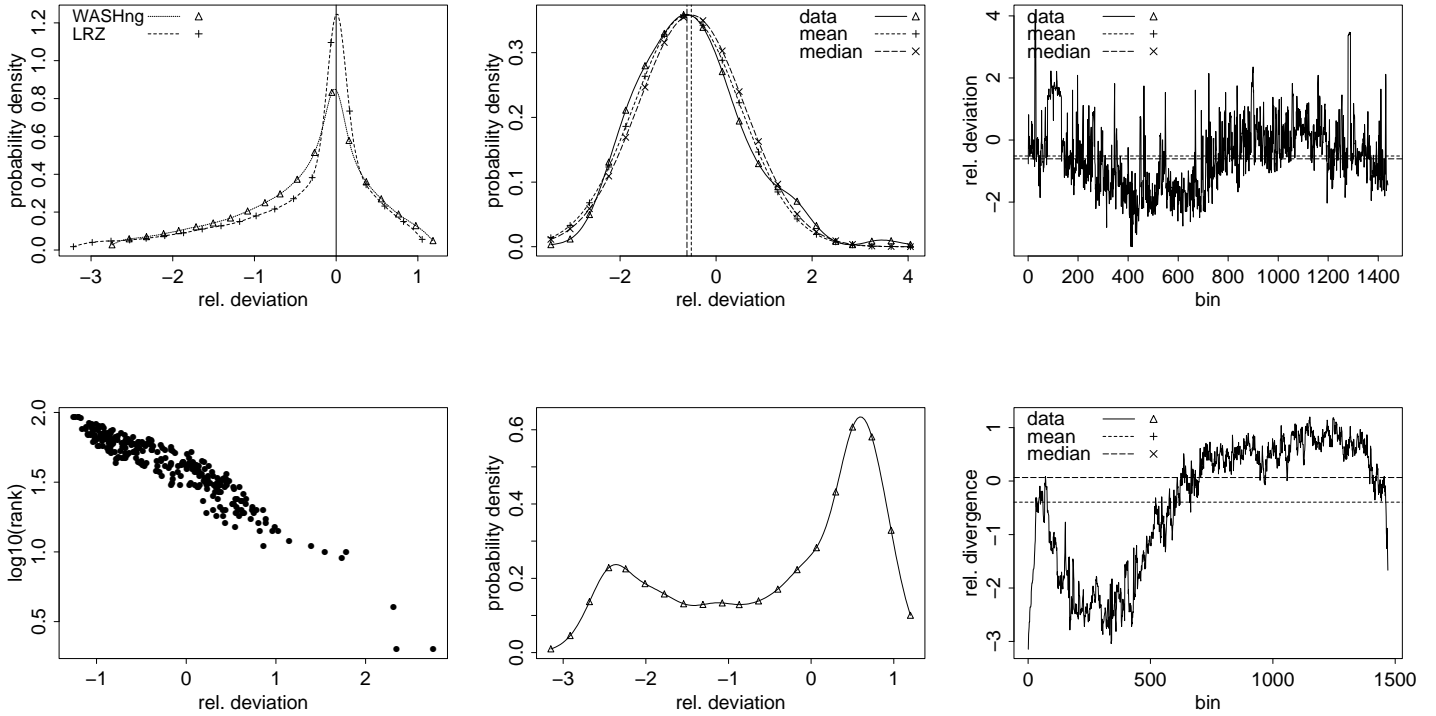
C. Flow Bandwidth Variability

So far our analysis has ignored the important aspect of bandwidth fluctuation and its impact on rank changes. As flow rates are heavy-tailed only large deviations will influence the ranking of top flows. Correspondingly, we do not care much about small additive deviations. Therefore we choose to consider a multiplicative metric, called *relative deviation* (or *reldev*). We compute this metric for all bins of each flow in the following manner:

$$\text{reldev} = \log_2 \left(\frac{\text{actual byte contribution}}{\text{mean byte contribution}} \right)$$

A *reldev* value 1 (2) stands for a deviation from the mean by a factor of 2 (4) and a value -1 (-2) represents a deviation by a factor of $1/2$ ($1/4$).

Figure 3(a) (top) shows, for both data sets, the probability density of the relative deviation for all flows with average



(a) Relative Deviation for 1000 Flows (top) and correlation of rel. deviation and rank change for an example flow (bottom)

(b) Probability density of rel. Deviation for two example destination prefix flows

(c) Rel. deviation as a function of time for two example destination prefix flows

Fig. 3. Relative deviation as metric for bandwidth fluctuation

rank 1–1000 that span at least 10 bins for 1 minute bins. The plot omits all values outside the 5% and 95% quantiles. Both curves have a similar form and show a pronounced spike at 0 while covering a range from -3 up to $+1$. (This corresponds to flow rates between $1/8$ -th of the average bandwidth use up to two times the average rate.) With increasing bin sizes (plot omitted) the curve for MWN widens slightly while it almost stays the same for the WASHng data set. The spike around 0 indicates that the bandwidth of most highly ranked flows does not diverge significantly from their average bandwidth. The broad base in both curves highlights that there are in fact deviations from the average flow rates, especially towards lower rates. Some of the larger absolute *reldev* values are due to short bandwidth spikes or bandwidth drop-offs while other are due to time-of-day effects for long lasting flows.

Figures 3(b) plot the probability density of *reldev* for two example flows both lasting for more than 1400 1 minute bins. The values range from -3 to $+4$. The distribution in Figure 3(b) (top) is close to a normal distribution with the same mean and standard deviation. This is not the case for the plot in Figure 3(b) (bottom) which shows a bimodal distribution. To understand the origin of this bimodality Figure 3(c) shows how the *reldev* behavior changes over time. (Note that this corresponds to a normalized plot of the bandwidth use over time). Figure 3(c) (bottom) which corresponds to the density

shown in Figure 3(b) (bottom) shows a clear time of day effect. For Figure 3(c) (top) the effect while discernible is not as pronounced. Both plots show that there are significant deviations in the *reldev* values throughout the lifetime of the flow and that the changes are not necessarily predictable or restricted to some part of the flows lifetime. For other bin sizes the same effects are visible.

Given the size of the bandwidth fluctuations (factors of 2–8) one can expect them to have an impact on the flow ranking. To verify this we plot the relative deviation against the rank of the flows. Yet as the relative deviations are logarithmic we also consider the logarithm of the ranks of the flow. Figure 3(a) (bottom) shows, for a typical flow, the corresponding scatterplot (*reldev* (x-axis) vs. decadic logarithm of flow rank (y-axis)) for 1 minute bins. There is one data point for each time bin. The plot shows a clear correlation between rank and *relative deviation*. This is a clear indication that bandwidth fluctuations have a significant influence on the flow ranking. The linear relation visible in the double logarithmic plot (*reldev* as well as rank) relates well to the fact that the bandwidth use is consistent with Zipf’s law. For a change in the top 10 ranks one needs a much larger bandwidth change than for a change in the ranks 900 to 1000. Such a correlation is not just apparent for this flow but seems to be present for all flows for all bin sizes.

IV. SUMMARY

A common assumption underlying much of the recent work that takes advantage of the heavy-tailed nature of traffic is that large traffic flows are persistent over time. In this paper we explore why this is not necessarily the case. While Zipf's law holds across time the ranking of the flows in terms of bandwidth contribution changes rather drastically. These changes can be decomposed into entry and exit processes as well as bandwidth fluctuations. We show, using an offline analysis, that the rank at which a flow enters or exits a ranking cannot be predicted or used as an indicator whether a flow is a large one or not. Furthermore, the bandwidth fluctuations if measured according to the relative deviations have a tightly centered empirical distribution about the flows average rate and are correlated to the rank changes. Since bandwidth changes over time are significant and correlated to rank changes every rank estimation techniques based on historical data has to account for them. It is dangerous to assume that flows that may have had a high rank once will stay highly ranked. A flows' average bandwidth and thus its average rank appears to be a better predictor. Therefore techniques that rely on online approximation of the average bandwidth appear promising for identifying subsets of flows that account for a significant fraction of the traffic.

ACKNOWLEDGMENTS

Many thanks go to the Leibnitz Rechenzentrum, München for access to the packet level traces. The analysis of the NetFlow traces is made possible by work supported by the National Science Foundation (NSF), cooperative agreement nos. ANI-0129677 (2002) and ANI-9807479 (1988) and the NLANR Measurement and Analysis Group (NLANR/MNA). We thank Holger Dreger, Balachander Krishnamurthy, and Walter Willinger for valuable suggestions regarding the analysis of the data.

REFERENCES

- [1] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of ethernet traffic (extended version)," *IEEE/ACM Trans. Networking*, vol. 2, 1994.
- [2] F. Brichet, J. Roberts, A. Simonian, and D. Veitch, "Heavy traffic analysis of a storage model with long-range dependent on/off sources," *Queueing Systems*, vol. 23, 1996.
- [3] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson, "Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level," *IEEE/ACM Trans. Networking*, vol. 5, 1997.
- [4] V. Paxson and S. Floyd, "Wide area traffic: The failure of poisson modeling," *IEEE/ACM Trans. Networking*, vol. 3, pp. 226–244, 1995.
- [5] M. E. Crovella and A. Bestavros, "Self-similarity in World Wide Web traffic: Evidence and possible causes," *IEEE/ACM Trans. Networking*, vol. 5, no. 6, pp. 835–846, 1997.
- [6] A. Feldmann, A. C. Gilbert, and W. Willinger, "Data networks as cascades: Investigating the multifractal nature of Internet WAN traffic," in *Proc. ACM SIGCOMM*, 1998.
- [7] K. Park and W. Willinger, eds., *Self-Similar NetworkTraffic and Performance Evaluation*. Wiley-Interscience, 2000.
- [8] G. Zipf, *Human Behavior and the Principle of Least Effort*. Addison-Wesley, 1949.
- [9] M. Mitzenmacher, "A brief history of generative models for power law and lognormal distributions.," *Internet Mathematics*, 2003.
- [10] L. Breslau, P. Cao, L. Fan, G. Philips, and S. Shenker, "Web caching and Zipf-like distributions: Evidence and implications," in *Proc. IEEE INFOCOM*, 1999.
- [11] D. N. Serpanos, G. Karakostas, and W. H. Wolf, "Effective caching of web objects using zipf's law," in *ICME*, 2000.
- [12] W. Fang and L. Peterson, "Inter-as traffic patterns and their implications," in *Proceedings of the 4th Global Internet Symposium*, December 1999.
- [13] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, "Deriving traffic demands for operational IP networks: Methodology and experience," in *Proc. ACM SIGCOMM*, 2000.
- [14] Y. Zhang, L. Breslau, V. Paxson, and S. Shenker, "On the characteristics and origins of internet flow rates," in *Proceedings of the IEEE/ACM SIGCOMM*, 2002.
- [15] J. Wallerich, H. Dreger, A. Feldmann, B. Krishnamurthy, and W. Willinger, "A methodology for studying persistency aspects of internet flows," *ACM Computer Communication Review*, 2005.
- [16] K. C. Claffy and N. Brownlee, "Understanding Internet traffic streams: Dragonflies and Tortoises.," *IEEE Communications*, 2002.
- [17] A. Feldmann, N. Kammenhuber, O. Maennel, B. Maggs, R. D. Prisco, and R. Sundaram, "A methodology for estimating interdomain Web traffic demands," in *Proc. ACM Measurement Conference*, 2004.
- [18] C. Estan and G. Varghese, "New directions in traffic measurement and accounting," in *SIGCOMM*, 2002.
- [19] S. Uhlig and O. Bonaventure, "Implications of interdomain traffic characteristics on traffic engineering," Tech. Rep. Infonet-TR-2001-08, University of Namur, Belgium, 2001.
- [20] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, "Bgp routing stability of popular destinations," in *Proc. ACM Measurement Workshop*, 2002.
- [21] S. Kandula and D. Katabi, "TeXCP: Responsive yet stable traffic engineering," in *Proc. ACM SIGCOMM*, 2005.
- [22] N. Ni and L. N. Bhuyan, "Fair scheduling and buffer management in internet routers," in *INFOCOM*, 2002.
- [23] L. Qiu, V. N. Padmanabhan, and G. M. Voelker, "On the placement of web server replicas," in *INFOCOM*, 2001.
- [24] K. Papagiannaki, N. Taft, and C. Diot, "Impact of flow dynamics on traffic engineering design principles," in *Proc. IEEE INFOCOM*, 2004.
- [25] K. C. Claffy, H.-W. Braun, and G. C. Polyzos, "A parameterizable methodology for Internet traffic flow profiling," *IEEE Journal on Selected Areas in Communications*, 1995.
- [26] S. Lin and N. McKeown, "A simulation study of IP switching," in *Proc. ACM SIGCOMM*, 1997.
- [27] Cisco Netflow.
<http://www.cisco.com/warp/public/732/Tech/nmp/netflow/index.shtml>.
- [28] "Abilene observatory."
<http://abilene.internet2.edu/observatory/>.