Impact of Flow Dynamics on Traffic Engineering Design Principles

Konstantina Papagiannaki†, Nina Taft‡, Christophe Diot‡
†: Sprint ATL, ‡: Intel Research
e-mail: dina@sprintlabs.com, nina.taft@intel.com, christophe.diot@intel.com

***Abstract*— A common traffic engineering design principle is to select a small set of flows, that account for a large fraction of the overall traffic, to be differentially treated inside the network so as to achieve a specific performance objective. In this paper we illustrate that one needs to be careful in implementing such an approach because there are tradeoffs to be addressed that arise due to traffic dynamics. We demonstrate that Internet flows are very volatile in terms of volume, and may substantially change the volume of traffic they transmit as time evolves. Currently proposed schemes for flow classification, although attractive due to their simplicity, face challenges due to this property of flows. Bandwidth volatility impacts the amount of load captured in a set of flows, which usually drops both significantly and quickly after flow classification is performed. Thus if the goal is to capture a large fraction of traffic consistently over time, flows will need to be reselected often. Our first contribution is in understanding the impact of flow volatility on the classification schemes employed in a traffic engineering context. Our second contribution is to propose a classification scheme that is capable of addressing the issues identified above by incorporating historical flow information. Using actual Internet data we demonstrate that our scheme outperforms previously proposed schemes, and reduces both the impact of flow volatility on the load captured by the selected set of flows and the required frequency for its reselection.**

I. Introduction

Aggregation of traffic according to source/destination network prefixes [10] has been proposed as an attractive network level abstraction useful in the engineering of large-scale IP networks [4]. At this level of abstraction, it has been shown that a small subset of the "flows" contributes to a large fraction of the volume [4]. Depending upon the context, these flows are sometimes referred to as heavy hitters or elephants.

Based on this initial observation there have been several proposals for the exploitation of this subset of the flows within a traffic engineering context that operates at the time scale of minutes or hours. The idea behind this traffic engineering principle is that if these large flows can somehow be identified, or picked out from among all flows traversing a given point, then they could be treated differently. One obvious application is load balancing [7] in which one could load balance a large portion of the traffic while only handling a small number of flows, and thus keeping only a small amount of state. There are other applications that could capitalize on this property of Internet traffic. For instance, the complexity of optimization algorithms for multipath routing (that tries to select equal cost multipaths for prefix-level flows) or route pinning through MPLS tunnels can be reduced if one only needs to do this for a limited set of flows [1], [8], [9], [11].

There have been several approaches proposed in the literature for the identification of heavy-hitter flows. For the benefit of simplicity, these approaches usually classify flows based on their behavior during a single time interval. The specific criterion used for the separation of the heavy-hitter flows out of the set of active flows is usually set in an arbitrary fashion. Typical examples include the identification of heavy-hitter flows as the top-N flows in the set, the flows that exceed a predefined bandwidth value B, or the heaviest flows that account for some fixed percentage X of the overall load.

By examining packet-level traces from an IP backbone, we show that destination network prefixes are volatile in terms of volume and this poses problems for such straightforward schemes when employed in a traffic engineering context. The set of flows that meet the specific flow classification criterion (e.g. account for 80% of the overall traffic or are the 100 highest bandwidth flows on the link) is changing frequently across time. Thus, if one wants to consistently achieve the flow classification target, flows would need to be often reselected. The difficulty with these approaches is that they focus on a flow's bandwidth at a single instant of time, and rely on ad hoc selection criteria. It is the volatility, or bandwidth fluctuations, that lead to frequent reclassification of flows over time. These fluctuations can come from two behaviors: i) short-lived bursts or drops in flow bandwidth, or ii) long-lived changes in the flow nature (e.g, long lasting shift from elephant to mouse, or termination of a flow).

This behavior makes it hard to identify flows that are both heavy and remain so for longer periods of time. It is preferable, however, in traffic engineering applications for flows classified as heavy flows to exhibit *persistence in time* so that frequent reclassification can be avoided. Reclassification can involve updating state information and can become a burdensome overhead if performed too frequently. In addition, it may affect the performance received by a flow if this flow is treated differently across distinct time intervals that are relatively close to one another.

To avoid reclassification of flows due to short-lived bursts or drops in their bandwidth we propose a new classification scheme that remains simple enough to be practical. Our scheme automatically detects a separation threshold that changes itself over time as a function of the flow statistics. It incorporates historical flow information and makes its decision based on a new per-flow metric that we call "latent heat". The latent heat scheme smooths historical flow bandwidth measurements, as well as the identified separation threshold, to minimize the number of reclassifications across time. We show that this new approach leads to a significantly more stable set of flows.

The flows selected as heavy-hitters at one point in time, however, may undergo significant changes later on that could manifest themselves as lower bandwidth values or even complete inactivity. We thus also need to determine how frequently flows need to be reselected so that they capture the targeted load. We show that if the desired goal is to capture a large volume of traffic (e.g., 80 or 90%) in the heavy-hitter class then the set of flows will need to be reselected at frequent intervals. Long-lived changes in the nature of heavy-hitter flows may result in significant fluctuations in the captured load. At the time of classification the carried load may meet the desired criterion but this load significantly drops as time evolves. We illustrate that if network designers could contend with capturing less load (e.g., 30% or 40%), then they would not need to reclassify as often because they are more likely to be able to detect persistent heavy hitters under such conditions.

In summary, in this work we illustrate that the notion of extracting heavy hitters is easier said than done. Our lessons learned indicate that network designers planning a traffic engineering application need to address two points: (i) how much load they really want or need to capture, and (ii) how often they are willing or able to reclassify. Our work highlights the trade-off between these two requirements.

The remainder of the paper is structured as follows. In Section II, we present the data set used to validate our heavy-hitter detection algorithm. Evidence on the bandwidth volatility of destination network prefix flows is presented in Section III. In Section IV we describe three different classification schemes that could be selected for the identification of heavy-hitters. In Section V, we demonstrate that single-instant classification schemes cannot identify a set of heavy-hitter flows that persists in time. Historical information is incorporated in our classification scheme in Section VI and evidence on the performance improvement is presented in Section VII. Limitations and recommendations for traffic engineering applications are also presented in Section VII. We conclude in Section VIII.

II. Experimental Environment

A. *Definition of a Flow*

A destination network prefix flow captures the traffic toward a specific set of IP addresses, as defined in the BGP routing table. A network prefix is the smallest routable entity in the Internet, and thus is an attractive network level abstraction for traffic engineering applications. At this traffic aggregation level, flows are limited in number and can be handled by changes in their next-hop IP address in the routing table [10].

Depending on the traffic observation point inside the network, we may see more or less traffic flowing toward any particular network prefix. For example, links inside a Point of Presence (PoP) that are close to the core of the network aggregate traffic from multiple routers. Therefore, they are likely to see more traffic toward specific destination prefixes than an access link inside the same PoP.

For the analysis presented throughout the paper, we have selected links that connect access routers to core routers inside the same PoP (OC-12 links). In this case, traffic is captured on its way to the core of the network. Our analysis has also been applied on OC-48 packet traces collected on links that interconnect core routers in different PoPs (OC-48 links), leading to findings similar to the ones presented throughout the paper. Given that the OC-48 traces never exceed 6 hours in duration, we will not use them in this paper.

For each flow, we measure its volume at fixed time intervals, and compute its bandwidth as the fraction of its volume over the duration of the interval. The measurement interval duration is set to 5 minutes. This choice is driven by operational practices: SNMP statistics are usually collected from routers every 5 minutes. At a 5 minute time scale, observations are also insensitive to protocol specific behavior that would impact flow bandwidth at timescales of seconds and milliseconds. In addition, traffic engineering applications that operate in large-scale IP backbone networks may not need to react at a finer time granularity. Even though our default time granularity is 5 minutes, we have also performed our analysis on flow measurements collected every 1 and 30 minutes. The results obtained were similar in nature to those presented.

B. *Measurement Data*

The data sets used in this paper come from packet traces collected inside Sprint's IP backbone network [5]. Optical splitters are used in conjunction with passive monitoring equipment to collect 44-byte headers from every IP packet traversing the monitored links. Monitoring equipment is deployed in four PoPs in the continental USA. Our algorithm has been systematically validated on 20 traces collected on OC-12 and OC-48 links, yielding similar results. To improve the readability of this paper, we present results from two OC-12 links (Table I). The utilization levels of the corresponding links are given in Figure 1 (the time displayed in the figures is always in PDT, i.e. UTC-8). We have chosen these two particular links because they correspond to the two longest packet traces at our disposal and thus allow for thorough analysis of the temporal behavior of network prefix flows. The two links are labeled according to their location within the continental U.S.

Trace	west coast	east coast
Start (PDT)	Jul 24 05:00:35 2001	Jul 24 05:00:34 2001
End (PDT)	Jul 28 23:42:55 2001	Jul 25 10:21:55 2001
#packets	1,677,983,111	1,678,055,791
Link Speed	OC-12 (622 Mbps)	OC-12 (622 Mbps)

TABLE I

Description of collected traces.

In parallel with the packet trace collection, we collect the BGP routing tables from route reflectors at the corresponding PoPs. The BGP tables are default-free and contain approximately 120K network prefixes. For each packet in a trace we perform a longest prefix match on the destination IP address. We then define the volume of the destination prefix as the sum of the payload of each packet destined to that particular network prefix.

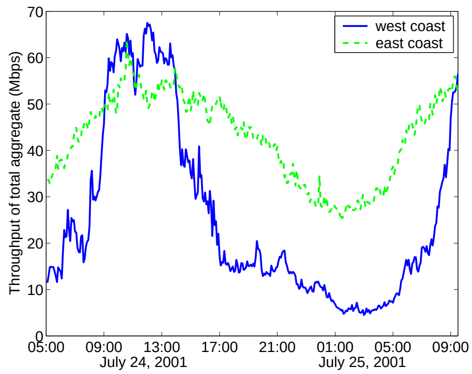

Fig. 1. Link utilization.

In our analysis we use a single BGP routing table that is collected at the beginning of the packet trace. Consequently, we cannot take into account changes in the BGP table that will occur during the trace. This simplification is not likely to affect our results since recent work has shown that the destination network prefixes that account for the majority of the load on a link rarely get affected by routing instabilities [6].

C. Initial Observations

We define a flow[1] to be *active* if it receives at least one packet during a measurement interval. Figure 2 shows the number of active prefixes in each measurement interval over a two day period. We find that in any given measurement interval, traffic travels toward approximately 10% of the network prefixes present in the BGP routing table.

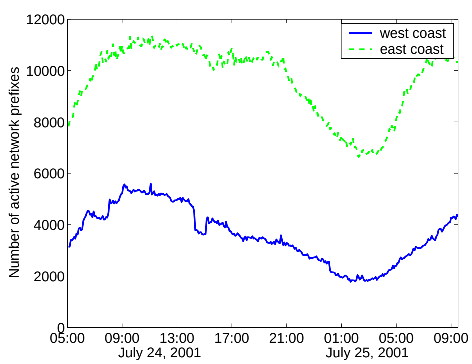

Fig. 2. Number of active network prefixes.

The cumulative distribution function (cdf) of the flow bandwidths is presented in Figure 3. Each curve corresponds to a different 5 minute interval within the first hour of the trace. The maximum number of active flows in each 5 minute interval for our two traces is in the order of a few thousands. Out of those the top 100 flows in each time interval systematically account for 50% to 70% of the total traffic, while the top 400 flows account for 75%-85% of the total traffic. Thus, we confirm that a small number of flows does account for the majority of the traffic in our trace.

[1] Throughout the paper we will use the terms "flow" and "prefix" interchangeably.

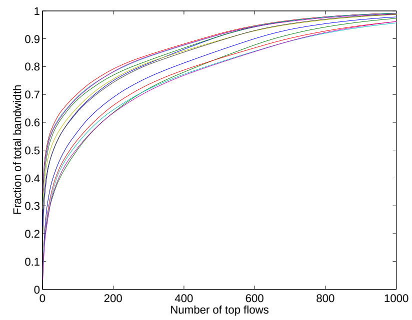

Fig. 3. Cumulative Distribution Function of flow bandwidths for the first 12 5-minute intervals of the west coast trace.

III. Indications of Volatility

Examining the set of flows in our data we found (not surprisingly) that all types of flows - in terms of volatility - exist. There are both heavy and light flows that are not volatile; i.e., heavy flows that exhibit steady behavior in that they remain heavy, and light flows that remain light for long periods of time. There are many flows (either heavy or light) that are prone to short-lived or transient bursts, and there are also flows that experience long-term changes. In this paper we do not address the reasons behind this volatility (addressed for 5-tuple flows in [12]) but rather its impact on traffic engineering.

In order to understand the prevalence of each type of flows and the extent of their volatility, we now look at our data from four different points of view. These perspectives serve as hints at the level of volatility of network prefix flows. First we compare the top 100 flows for all pairs of sequential intervals during the first day of both traces. We observe that the set of flows constituting the top 100 changes significantly from one time interval to the next. The minimum number of flows that change between two consecutive intervals is 23 and the maximum reaches 55 (i.e. half of the flows in the top 100 flows at one interval are replaced by new flows in the next interval).

To identify the reason behind this phenomenon, we collect flow measurements for the first hour of the west-coast trace. For each flow we measure its average bandwidth across the first hour of the trace as well as its coefficient of variation (e.g. the fraction of the standard deviation of the flow bandwidth divided by its average). In Figure 4 we present the relationship between the two metrics for all the active flows.

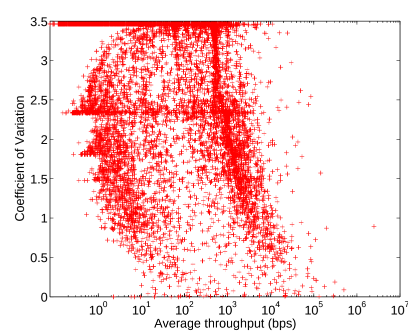

Fig. 4. Relation between the coefficient of variation of flow bandwidth against average bandwidth for the 1st hour of the west coast trace.

Figure 4 shows that there is no clear correlation between the mean and the coefficient of variation of the bandwidth of a network prefix flow. Flows of any size (in terms of average throughput) can have either a small or a large coefficient of variation. Since the center and top portion of the graph are darker, it means that most flows do experience variation. Any (not "all") flow may experience fluctuations around their average bandwidth that reach up to three times its value. Consequently, small-volume flows may experience bursts that raise their throughput to significant levels while large-volume flows may experience drops that significantly reduce their overall volume. Since there are many flows in the intermediate throughput range and they too have a very large variance, it means any of them could become a heavy hitter at some point or could become an insignificant flow at some other point.

We next ask the following question. Suppose we were to select the set of flows constituting 80% of the load at some point in time. If we track those flows without changing the set of flows, what fraction of the total load would they constitute a few hours later? As an example, we select the highest bandwidth flows that account for 80% of the traffic at 9:00 am on the west coast trace. We then track the volume accounted for by these flows for the remainder of the day. As can be seen from Figure 5 when heavy flows are initially selected their corresponding cumulative load is indeed 80%. After one hour has elapsed, however, their overall load has dropped to 60%, and it reaches 45% at 11am. Similar results hold when the classification of the flows takes place at other times during the day.

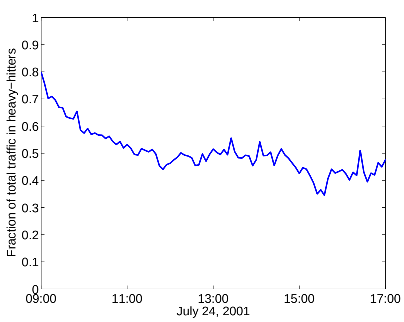

Fig. 5. Evolution of fraction of overall traffic in heavy-hitters initially selected to account for 80% of the load.

This phenomenon may be due to flows reducing their overall bandwidth or simply disappearing from inside the network, while flows that have increased their bandwidth are not taken into account since the set of heavy-hitters is static. To evaluate the effect of flows becoming inactive with time we measure the lifetime of a flow as the number of intervals that a flow is active between periods of inactivity. We present our results in Figure 6. We notice that even though there is a significant number of flows that last longer than 3 hours, 40% of the flows will end within 1 hour.

From these views, we can draw the following two main conclusions. First, there are many flows that span a large range of volume values and can become candidate heavy flows at some point in time. Furthermore, heavy flows may not be heavy for long periods of time and thus reassessment of which flows belong in the heavy class may be needed frequently. Second, this

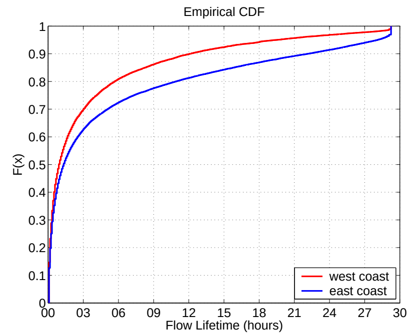

Fig. 6. Lifetime of a flow for the west and east coast trace.

can make the job of identifying a set of heavy hitters challenging if the goal is to consistently capture a large portion of the load but simultaneously avoid reclassifying traffic often.

In the next section, we evaluate three simple classification schemes for the identification of heavy-hitter flows in Internet traffic. We show that simple schemes that base such a decision on the instantaneous properties of Internet flows are not likely to produce a good set of flows that traffic engineering applications could use.

IV. Single-Instant Classification Schemes

Classification schemes proposed in the literature usually classify flows according to their "instantaneous" behavior, e.g. their properties during a single time interval. We call these schemes "single-instant classification schemes".

Let i denote the index of a network prefix flow in a BGP routing table. Let τ denote the length of the time interval over which measurements are taken[2]. Time is discretized into these intervals, and n is the index of time intervals. We define $X_i(n)$ to be the average bandwidth of the traffic destined to a particular network prefix i during the n^{th} time slot of length τ. Whenever flows are separated into two classes according to a threshold value, this value is denoted as $v(n)$ (bps). We use $C_1(n)$ to represent the set of flows contributing to the higher fraction of traffic at interval n (we call these flows the "heavy-hitters") and $C_2(n)$ for all other flows. That is

$$if\ X_i(n) \geq v(n),\ i \in C_1(n),\ else\ i \in C_2.$$

Notice that the set of flows identified as "heavy-hitters" depends on the time interval n and the respective value of $v(n)$. Therefore, sets C_1 and C_2 are functions of n.

A. Existing Schemes

In this section we describe two previously proposed classification schemes. Both these schemes require the setting of an input parameter that guides the selection of heavy hitter flows. We call the first scheme the "constant load" scheme. Given a fraction l, this scheme selects heavy hitters to be the highest bandwidth flows that account for a fraction l of the overall load on the link. The second scheme does not target a specific load

[2] In this work τ is set to 5 minutes but could take on different values depending on the requirements of the traffic engineering application.

in the heavy-hitter class, but rather a specific number of flows in it. We call this scheme the "top-N" scheme, which isolates as heavy-hitters the N highest bandwidth flows on the link. A third scheme that is frequently suggested in the literature selects heavy-hitters to be these flows that exceed a pre-specified bandwidth value B [3]. We do not evaluate this approach since it assumes pre-existing knowledge for the range of bandwidth values achieved by flows.

Notice that the first two schemes suffer from two obvious weaknesses: 1) they require the setting of parameters l and N, and 2) they are sensitive to outliers. For the "constant load" scheme there is always the danger that a flow that bursts at high bandwidths may account for a large fraction of the total volume and thus mask out multiple other high-bandwidth flows that may exhibit greater persistence in time. For the "top-N" scheme, flows bursting for brief periods of time to high values of bandwidths can easily displace other flows from the top N despite their being more persistent in time.

To address the above two weaknesses we propose a new scheme which we call the *aest* scheme. This scheme does not require any input parameter and features properties that make it less sensitive to outliers.

B. Automatic Detection Scheme: aest

The method we propose relies on our observation that the flow bandwidth distribution is heavy tailed. We analyze the flow bandwidth distributions collected for each 5 minute interval in our traces and use the *aest* test [2], formally defined later in this section, to evaluate whether they have heavy-tail properties. Our results reveal that all collected distributions are characterized by a scaling exponent α between 1.03 and 1.2. Consequently, we can capitalize on this property and isolate the flows that fall in the tail of the bandwidth distribution as the heavy-hitters during a time interval. According to this method we can set $v(n)$ as the cutoff point in the distribution after which power law properties can be witnessed. This approach has several interesting properties. First, $v(n)$ is not a constant value. Instead, it follows the dynamics of the traffic and adapts to its changes due to time of day behavior, for example. Second, $v(n)$ can be automatically computed. Third, $v(n)$ is not sensitive to outliers. One high-bandwidth flow is not likely to mask another high-bandwidth flow, since they will both fall at the tail of the flow bandwidth distribution.

The *aest* method identifies the portion of a distribution that exhibits power-law behavior. The key idea is that the shape of the tail determines the scaling properties of the dataset when it is aggregated. The aggregation of a dataset of N observation X_i, $i = 1, \cdots, N$, is defined as the process of summing non-overlapping blocks of observations of size m: $X_i^{(m)} = \sum_{j=(i-1)m+1}^{im} X_j$. By observing the distributional properties of $X^{(m)}$ one can infer where in the tail power-law behavior begins. We use this method to identify the first point in the flow bandwidth distribution after which power law properties can be witnessed.

We present the graphical output of the *aest* tool in Figure 7. This figure presents the complementary cumulative density function (CCDF) for the aggregates of the flow bandwidth distribution collected in a single time interval n. The segment of

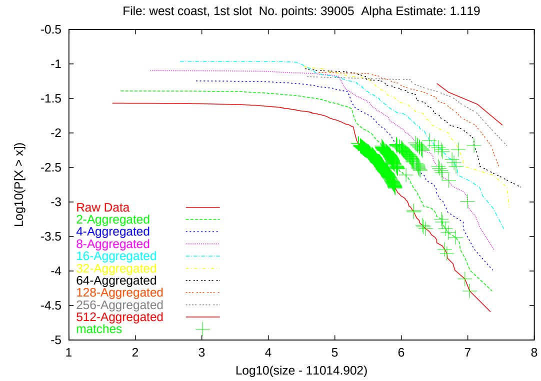

Fig. 7. Graphical output of the *aest* tool. Presents the complementary cumulative density function of the number of bytes sent by flows during the first 5 minute interval of the west coast trace.

the tail over which heavy-tailed behavior appears to be present is denoted with "plus" signs. We select the minimum value of bandwidth after which heavy-tailedness is confirmed as the cut-off point in the original flow bandwidth distribution. This value in Figure 7 is 5.35, that corresponds to $10^{5.35}$ bytes sent over that 5 minute interval. We use this particular bandwidth value as the threshold value $v(n)$. If we define $\bar{F}(x)$ as the complementary cumulative distribution function of the flow bandwidths $X_i(n)$ at interval n, i.e. $\bar{F}(x)$ = 1 - F(x) = P[X>x], then

$$v(n) = \min(x : \frac{d \log \bar{F}(x)}{d \log x} \sim -\alpha)$$

The *aest* methodology has the advantages of being fast, non-parametric and easy to apply. It has been shown to be relatively accurate on synthetic datasets [2]. More importantly, there has been evidence that the scaling estimator α, as calculated by *aest*, appears to increase in accuracy as the size of the dataset grows. Given that our datasets feature at least two thousand measurements (i.e. active flows per time interval), *aest* will provide us with reasonably accurate values for the cutoff points in the heavy tail flow bandwidth distribution measured in each time interval.

V. RESULTS ON SINGLE INSTANT CLASSIFICATION SCHEMES

In most applications of classification each data object inherently belongs to some particular class. Our situation is different in that the flows do not naturally belong to a particular class. Consequently, traditional evaluation techniques based on misclassification errors cannot be used. Given that we are addressing the traffic engineering principle of isolating heavy hitters and the feasibility issues of doing so consistently over time, we focus on temporal persistence. We, thus, evaluate the alternative schemes with respect to their capability of providing a set of flows that does not significantly differ from one interval to the next.

The length of time that a flow in C_1 will remain in C_1 is both a function of the flow itself and of the classification. It is a function of the classification in the sense that a particular high-bandwidth flow will remain in C_1 as long as the continually adjusting threshold stays lower that its average bandwidth. The classification scheme that produces a set of flows that maintain their classification for longer periods of time is more attractive for our purposes.

To define a performance metric to capture the persistence of the classification in time, we proceed as follows. The proposed classification scheme induces the following underlying two-state process on each flow. Let $I_i(n) = 1$ if $i \in C_1(n)$, and $I_i(n) = 0$ if $i \in C_2(n)$. At each classification time interval, the process either transitions to the other state or stays in the same state. To capture the persistence of the classification in time, for each flow i, 1) we count how many times it transitions to C_1, and 2) we measure the number of sequential time intervals the flow spends in C_1 for each stay. We then compute the average holding time of a flow in class C_1 as the average number of sequential intervals a flow has been in C_1.

To compute the average holding time of a flow according to the above description we need to select the time interval when we start accounting for flow transitions (denoted as n_1) and for how long ($n_2 - n_1$, where n_2 denotes the end of the period). If the period over which we compute the average holding time of a flow is greater than 12 hours, then the flow behavior is sensitive to the diurnal cycle of the traffic (as shown in Figure 1). In this case, small holding times are not due to the inefficiency of the classification scheme to produce a persistent set of flows in C_1, but due to the significant decrease of $X_i(n)$ due to time of day fluctuations. To be able to evaluate the classification scheme in terms of time persistence we thus need to identify a period of time when flows are less likely to be affected by the diurnal cycle and dramatic transitions in their behavior like before and after the working hours. The 5-hour peak period between 9 am and 2 pm is characterized by a stable number of flows and stable link utilization for both traces. In addition, it features the largest number of flows that contribute the greatest amount of load across the entire day. Therefore it lends itself as the most appropriate period when classification persistence can be tested. Consequently, we set n_1 and n_2 to the time intervals that correspond to 9 am and 2 pm respectively.

We present statistics on the average holding time in C_1 in Table II for both traces under the three alternative single-instant classification schemes[3]. Under the two previously proposed schemes average holding times in C_1 are between 15 and 20 minutes. More than 80% of the flows experience an average holding time in C_1 of less than 20 minutes for both traces. Using the *aest* classification scheme the holding time in C_1 slightly improves but is still limited to less than 1 hour for 80% of the flows for the west coast trace, while it appears insignificant for the east coast trace. Therefore, the aest scheme is not capable of achieving persistence on its own. In the next section, we introduce a new scheme that is capable of addressing the limitations of aest to achieve a higher degree of persistence in the set of the heavy-hitter flows.

[3]In this table we report on the performance of the constant load scheme when it targets 70% of the total load in class C_1. Our results across different constant load schemes showed that the average holding time in C_1 never exceeds 40 minutes, achieved by the 0.9-constant load scheme for the west coast trace.

West Coast			
	top-100	0.7-constant load	*aest*
min.	1	1	1
20%	1	1	1
50%	1.2	1.6	2
avg.	3.47	4.7	8.63
80%	2.75	4	12
max.	60	60	60
East Coast			
	top-100	0.7-constant load	*aest*
min.	1	1	1
20%	1	1	1
50%	1	1.5	1.5
avg.	3.52	3.7	3.72
80%	3	3	3.5
max.	60	60	60

TABLE II

STATISTICS FOR THE HOLDING TIMES IN C_1 (IN 5-MINUTE SLOTS) UNDER 3 SINGLE-INSTANT CLASSIFICATION SCHEMES.

Frequent changes in a flow's class may be due to two reasons. Firstly, short-lived bursts or drops in the flow's bandwidth. In Figure 4 we showed that flows may significantly deviate from their average bandwidth value throughout time. Secondly, a flow may get reclassified because of genuine long-lived changes in its nature. To differentiate between these two phenomena, it would be advantageous if our scheme could incorporate historical flow behavior in the classification process. Smoothing operations on the flow bandwidth measurements as well as the computed separation threshold $v(n)$ could limit the effect of short-lived bursts or drops in the flow bandwidth and lead to higher levels of persistence for set C_1.

VI. LATENT HEAT CLASSIFICATION SCHEME

We introduce the notion of the "latent heat" to allow a flow to maintain its heavy-hitter classification despite brief transitions across the threshold $v(n)$. With the latent heat we allow large flows to experience short *transient* periods in which their volume drops. Similarly, flows in C_2 are not allowed into C_1 due to *transient* bursts above $v(n)$.

A. Methodology

In order to detect the points in time when flows transition from one class to the other, we introduce a new per flow metric, namely the "latent heat"[4]. A flow is assumed to accumulate energy when it is above $v(n)$ and to lose energy when it is below $v(n)$. We define $C_1'(n)$ to be the set of flows whose latent heat exceeds a threshold value $thres$. Transition of a flow i from $C_1'(n)$ to $C_2'(n+1)$ is triggered when the latent heat of a flow falls below $thres$. Transition of a flow i from $C_2'(n)$ to $C_1'(n+1)$ is triggered when the latent heat of a flow increases beyond $thres$.

[4]The term "latent heat" comes from thermodynamics and defines the heat energy required to change a substance from one state to another.

The latent heat scheme can be conceived as a three-step scheme. Firstly, we compute the value for $v(n)$ using the *aest* approach presented in Section IV. Then, we measure the "heat" of each flow as the difference between the flow's throughput $X_i(n)$ and $v(n)$, which we denote by $d_i(n)$. Lastly, we apply an exponential weighted moving average (EWMA) filter on the difference[5] $d_i(n)$ and compare it to a prespecified threshold value $thres$. The smoothing constant is denoted as γ, $0 < \gamma \leq 1$. We present the formulation of the scheme in the following equations.

$$d_i(n) = X_i(n) - v(n) \tag{1}$$

$$LH_i(n) = \begin{cases} (1-\gamma)\, LH_i(n-1) + \gamma\, d_i(n-1) & \text{if } n > 1 \\ 0 & \text{if } n = 1 \end{cases} \tag{2}$$

$$if\ LH_i(n) > thres,\ i \in C_1'(n),\ else\ i \in C_2'(n). \tag{3}$$

In essence, the latent heat could be any function of the difference between bandwidth and $v(n)$ over a number of intervals. We chose to use an EWMA filter because it requires minimal historical flow measurements. Our formulation of the latent heat metric assigns equal weight to a flow's drop and burst. In case a network engineering application values a flow's rate drop more than a flow's burst, then Eq. 2 can be adjusted accordingly to reflect this requirement.

Notice that the set of flows in $C_1'(n)$ is not the same as the set of flows in $C_1(n)$. The first set contains all the flows that exceed $v(n)$, while set $C_1'(n)$ contains all the flows that have a latent heat greater than $thres$. In terms of their properties, class $C_1(n)$ contains all the high-volume flows during a single time interval, while $C_1'(n)$ contains all the flows that have acted as a high-volume flow in the past and are more likely to continue being high-volume in the future. We call the flows in C_1 heavy-hitters and the flows in C_1' *persistent heavy-hitters*.

We illustrate the notion of the "latent heat" in Figure 8. We show the bandwidth measured for one particular destination network prefix on the west coast trace along with the computed threshold $v(n)$ (automatically detected with the *aest* approach), the flow's latent heat and $thres$ (equal to zero in this example). Until 8 am the flow's bandwidth is below $v(n)$, causing its latent heat to slowly decrease. At approximately 8 am the flow's rate increases significantly. It starts accumulating energy and at 8:30 am the flow's latent heat becomes positive. Since $thres$ is equal to 0 this is the first point in time when this flow is placed in C_1'. Until 11 am this flow achieves significantly high bandwidth rates, and therefore its latent heat continues to increase. At 11 am the flow's bandwidth falls below $v(n)$ and its latent heat starts to decrease rapidly (the slope of decrease is much steeper than during the period between 5 am and 8 am due to the significant change in bandwidth). At 12 am when the latent heat of the flow has approached zero, the flow resumes its high rates increasing its energy. After 2 pm the flow's bandwidth approximates zero, and then its latent heat falls below zero until at least 9 pm.

[5] Note that smoothing the difference between $X_i(n)$ and $v(n)$ is equivalent to the difference between the smoothed flow bandwidth $X_i(n)$ and the smoothed separation threshold $v(n)$.

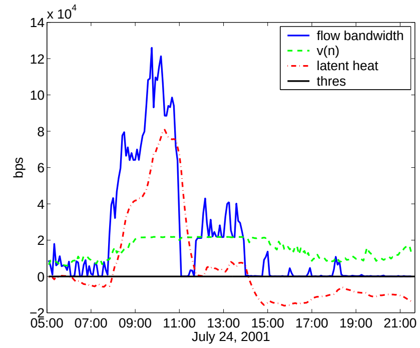

Fig. 8. Illustration of the idea behind the "latent heat" metric.

If the classification methodology does not take into account transient changes in state then this flow would be placed in the low class for the drop after 11am. In fact, under the *aest* scheme, the flow would have been reclassified 6 times between 5 am and 2 pm. With the latent heat metric reclassification occurs only twice, once at approximately 9 am and once at 2 pm. This behavior is beneficial in a traffic engineering context since the application would have to be consistently treated as a high-rate flow for the period between 8am and 2pm.

B. *Impact of γ and $thres$*

Parameter γ in Eq. (2) determines the weight of the current flow behavior on the value of the EWMA and the number of historical measurements taken into account. Lower values of γ assign less weight on current flow measurements and more weight on historical behavior. The average number of past intervals K that are accounted for in the EWMA relates to the value of γ according to the following equation (see Appendix for proof).

$$\gamma = \frac{2}{K+1} \tag{4}$$

If a flow has historically been performing as a heavy-hitter, then under small values of γ this flow is likely to continue being classified as a heavy-hitter despite brief changes in its behavior. Similarly, under large values of γ, it will take longer for a flow to be admitted to C_1'. Parameter $thres$ can further limit the number of flows in C_1' targeting those flows that have accumulated larger amounts of latent heat. Setting $thres$ equal to 0 allows any flow that experiences brief periods of positive latent heat to be placed into C_1'.

To investigate the combined impact of the two parameters on the set of flows classified as persistent heavy-hitters we present the way they relate to the time persistence of class C_1'. Results are presented in Figure 9. An increase in the value of parameter $thres$ is accompanied by a decrease in the average holding time of flows in C_1'. Flows need to possess much greater values of latent heat to continue being classified in C_1', and thus stay in that class for smaller periods of time. Similarly, an increase in the value of γ leads to smaller average holding times in C_1'. Greater values of γ incorporate less past flow behavior and therefore short-lived changes in a flow's bandwidth are more likely to lead to a flow's reclassification.

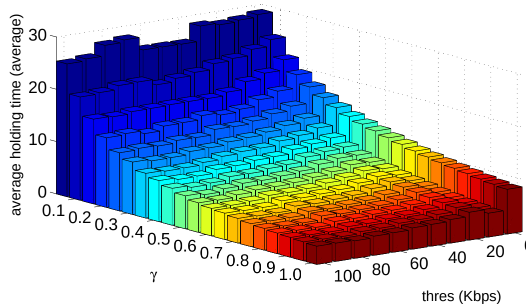

Fig. 9. Average of flow average holding times vs. γ and $thres$ (west coast).

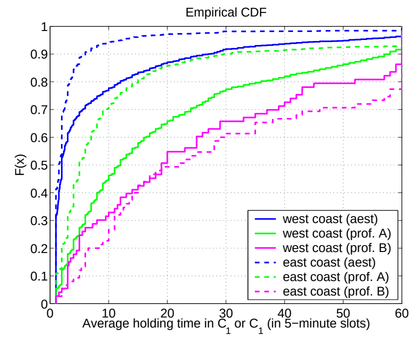

Fig. 10. Comparison of the *aest* approach with the latent heat under $\gamma = 0.2$, $thres = 0$ (Prof. A), and $\gamma = 0.05$, $thres = 100Kbps$ (Prof. B).

For the remainder of this section, and due to space limitations, we focus our attention to two latent heat schemes alone. They represent two configurations that could be selected in a traffic engineering context. The first scheme (Profile A) is characterized by $\gamma = 0.2$ and $thres = 0$ Kbps and the second scheme (Profile B) by $\gamma = 0.05$ and $thres = 100$ Kbps. Both schemes account for historical flow behavior and thus lead to more persistent sets of heavy-hitter flows. The first scheme accounts for approximately 9 intervals (e.g. 45 minutes) of historical information and accepts in C_1' any flow with a latent heat greater than 0. On the other hand, the second scheme accounts for 39 intervals, on average, of historical information (e.g. approximately 3 hours) and restricts the set of heavy-hitter flows to only those flows that have a latent heat greater than 100 Kbps. Thus profile B will accept fewer flows into the heavy hitter class, will capture less overall load, but will select flows with greater persistence than profile A.

VII. Results

A. Flow classification time persistence

In order to demonstrate the benefits of our proposed approach we compare the persistence of set C_1, produced by *aest*, with the persistence of set C_1', produced with the latent heat schemes, in Table III. For the latent heat scenario we use configurations A and B. We show that single-instant classification schemes, like the *aest* approach, lead to an average lifetime in C_1' for 15 to 40 minutes (Table III). The average holding time achieved by the latent heat scheme under profiles A and B is greater than 1 hour. In Figure 10 we present the CDF for the average holding time in C_1' for the *aest* and the latent heat schemes. Figure 10 provides clear evidence on the performance improvement, in terms of time persistence for set C_1', achieved by the latent-heat scheme. Under the *aest* approach more than 90% of the flows experience an average holding time in C_1' of less than 1 hour. Under the latent heat scheme, in three out of the four cases, less than 50% of the flows experience similar average holding times. The one exception applies to the east coast trace for profile A, which still succeds to increase the number of flows experiencing average holding times more than 1 hour from 5% up to 25%.

	West Coast			East Coast		
	aest	Pr. A	Pr. B	*aest*	Pr. A	Pr. B
min.	1	1	1	1	1	1
20%	1	4	3	1	2	3
50%	2	12	19	1.5	5	22
avg.	8.63	20	17.3	3.72	12	16.2
80%	12	36	28.5	3.5	15	24.5
max.	60	60	60	60	60	60

TABLE III

Statistics for the holding times in C_1' under the *aest* and the latent heat approach (Prof. A: $\gamma = 0.2$, $thres = 0$, Prof. B: $\gamma = 0.05$, $thres = 100Kbps$) (in 5-minute slots).

B. Frequency of Reclassification

In the previous section we evaluated the capabilities of three different classification schemes with respect to the persistence of the set of the heavy-hitter flows they produce. For the evaluation of the time persistence of the produced flow classification we compared the set of flows classified as heavy-hitters across multiple time intervals. Therefore, the obtained results reflect how the flow classification would behave if one were to reselect the heavy-hitter flows at every time interval. In this section, we look into the impact of selecting the heavy-hitters at one point in time and treating them as heavy-hitters for a greater period of time. To evaluate this effect we classify flows at one instant in time and look into the amount of traffic they account for as time elapses.

Given that results may differ depending on the time of day when the classification takes place, we evaluate the persistence of the heavy-hitter load as follows. We apply our classification schemes at the beginning of each hour in our packet traces and track the fraction of total traffic that the selected flows carry through the next hour.

1) Constant load schemes: Under the constant load schemes the classification process attempts to capture a specific fraction of the load in the heavy-hitters. In Figure 11 we present how much the achieved captured load diverts from the targeted load up to one hour after the selection of the heavy-hitters for different l-constant load schemes. Typical loss of load for schemes targeting more than 60% of the traffic is approximately 25%, i.e. the scheme targets 60% and it usually

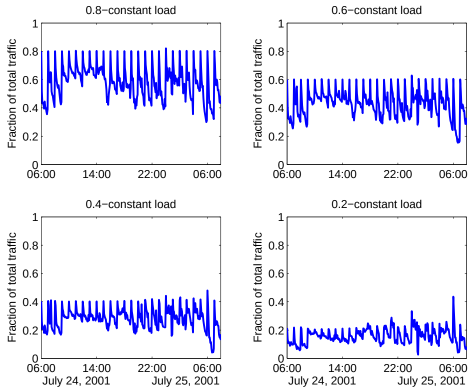

Fig. 11. Hourly loss in captured load under constant load schemes.

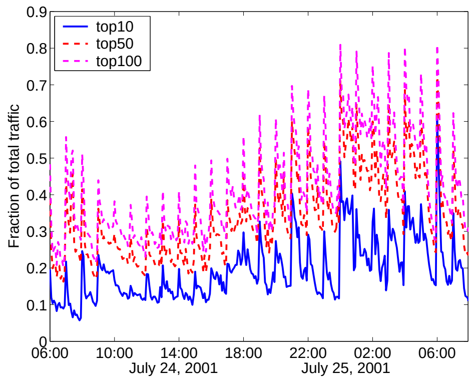

Fig. 12. Hourly loss in captured load under top-N schemes.

reaches 35% if the heavy-hitters are not reselected for one hour. When the targeted load is greater than 30% losses in load are typically on the order of 10-20%. Lastly, when the targeted load is below 30% then the fluctuations in load are more contained.

2) Top-N schemes: To further investigate this phenomenon, we perform the same analysis on the top-N schemes, whose goal is to capture the N highest bandwidth flows on the link. We select these flows at the top of each hour and track their overall load for the remainder of each hour. Our results are presented in Figure 12.

We see the same kind of behavior as before, namely that the fraction of load captured drops significantly within the hour and quickly (within the first portion of the hour). This is true for all three values of N we considered (and has been tested for values of N up to 1000). We also observe in Figure 12 that the greater the number of top flows one selects as heavy-hitters the larger the fluctuations in the heavy-hitter load with respect to the overall link throughput. Selection of the heavy-hitters as the 10 highest bandwidth flows on the link leads to smaller fluctuation in the overall load during the peak hours but significantly increases after 6pm. Average fluctuations for the top-10 scheme are on the order of 10% while for the top-100 scheme they reach up to 50%. This implies that the more flows one selects as heavy-hitters the more frequently they should be reselected. Hence network designers would do better in terms of isolating a consistent amount of load if they shift their goal from "capturing the majority of the load" to capturing smaller fractions of it.

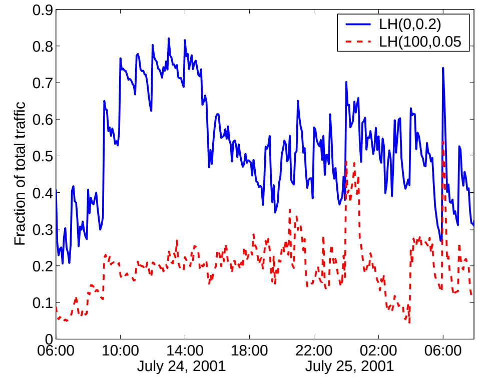

Fig. 13. Hourly loss in captured load under latent heat schemes.

3) Latent heat schemes: Lastly, we look into the performance of the latent heat scheme. Figure 13 shows that fluctuations in the load of class C_1' are significantly smoother under the latent heat schemes if reselection of flows were to take place once an hour. The first profile LH(0,0.2) is capable of capturing approximately 80% of the traffic during the peak hours with fluctuations of less than 10% (within approximately 1000 flows). On the other hand, the second profile LH(100,0.05) almost consistently succeeds in placing 20% of the overall traffic in less than 50 heavy-hitter flows. Given the limited fluctuations of the load under the second profile reselection of flows could be performed less frequently (e.g. every 2 hours). This scenario reiterates the point that by capturing less load (20% in this case), network designers can reclassify less often while simultaneously maintaining a more consistent fraction of the load.

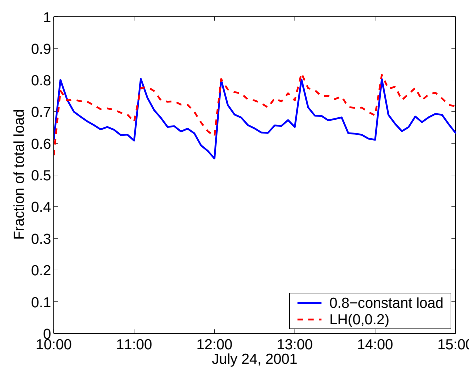

Fig. 14. Hourly loss in captured load for 0.8-constant load and LH(0,0.2) schemes across the busy period.

We notice that despite the fact that the LH(0,0.2) scheme does not target a specific load in C_1' it succeds in capturing approximately 80% of the total traffic for the peak hours of the trace. Nevertheless, it does so in a way that the fluctuations in the captured load are much smoother and using a more persistent set of flows across time. To better illustrate the first point we present a detail of Figure 13 focusing on the peak hours of the trace in Figure 14. The fluctuations in the captured load under the latent heat scheme are much smoother and less abrupt

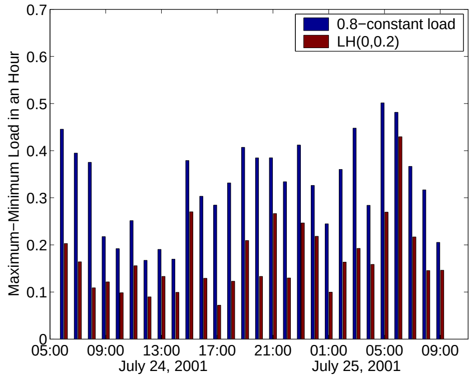

Fig. 15. Maximum loss in captured load for 0.8-constant load and LH(0,0.2) schemes across time.

than under the 0.8-constant load scheme. This load is offered by approximately 1000 flows. Given that these flows have an average holding time in C_1' greater than 1.5 hours (Table III) the set of flows selected across time will also be more consistent across time. To better quantify the loss in the captured load between flow reselections for the 0.8-constant load and LH(0,0.2) schemes we compute the difference between the maximum and minimum captured load across each hour in the trace. We present our results in Figure 15. The latent heat scheme always leads to smaller loss in the captured load thus smoothing out the effect of flow volatility on the accounted traffic load. Therefore, the latent heat schemes are beneficial both in terms of consistency in the captured load as well as persistence in the set of flows classified as heavy-hitters.

In summary, traffic engineering decisions applied to flows defined at the traffic aggregation of a destination network prefix should anticipate changes in the nature of the flows selected as heavy-hitters. If the goal is to capture a large (e.g., 70 or 80%) fraction of the traffic load, then flows will have to be reselected frequently, since they are unlikely to maintain the same load as time evolves. In essence, to affect a large amount of traffic one needs to select a large number of flows. Given that these flows have different features with respect to their lifetime and bandwidth volatility, the more flows one accounts for, the more likely it is that the observed behavior will significantly change in the future.

We have shown that application of single instant classification schemes for the selection of a small set of flows that will be exploited in a traffic engineering context is not sufficient. Significant changes in the load of the selected flows introduces large fluctuations in the amount of the captured traffic and flows need to be frequently reselected to lead to the desirable goal. Accounting for historical information is capable of filtering out short-lived changes in the bandwidth of the selected heavy-hitter flows producing a more stable set of flows to be exploited in a traffic engineering context. In addition, the latent heat scheme leads to less volatility in the overall load accounted by heavy-hitters. Nevertheless, long-term changes in the nature of heavy-hitters will still require flows to be reselected across time.

C. *Time of Day Influence*

From the previous section we saw that fluctuations in the load captured by heavy-hitters are usually more contained during the peak hours of the day. Thus, there is an indication that the persistence of heavy-hitters may be related to the time of day when they are selected. In fact, traffic in the Internet is known to follow a diurnal cycle, as also shown in Figure 1. In this section, we address the relation between the persistence of a heavy-hitter in time against the time of day when flows are classified.

To address this issue we make use of the first two days in the packet trace for the west coast link, i.e. Jul 24, 2001 05:00 until Jul 26, 2001 05:00. For each hour h after the beginning of the trace we collect the set of flows in C_1', according to LH profile A ($\gamma = 0.2, thres = 0Kbps$). For each flow $f \in C_1'(n)$ we count the number of sequential intervals in the future that this flow remains in C_1', which we call the future lifetime of flow f in C_1'. We report the latent heat of every flow f along with its future lifetime in C_1'. We then look into the median and average value of the future lifetime of a flow in C_1' according to its latent heat at time interval n for the period between 12 am on July 25, 2001 and 12 am on July 26, 2001. Notice, that the future lifetime of each flow in C_1' will be upper bounded by the ending time 05:00 on Jul 26, 2001.

Figure 16 presents the relation between the average future lifetime of a heavy-hitter at time n when its latent heat exceeds different values as presented on the x-axis. Results for the median value of the future lifetime of a persistent heavy-hitter are shown in Figure 17. There are a few interesting observations one can make on Figure 16 and Figure 17. First, if a flow at time interval n has a large latent heat $LH_i(n)$ then this flow is likely to remain in C_1' for longer periods of time than a flow with a small latent heat. This statement holds for any n and results start to deteriorate when $thres > 80$ Kbps. Therefore, the latent heat of a flow is a good predictor of its capacity to remain in C_1' in the future. Second, flows are more likely to stay in C_1' during the night rather than during the day. Flows in C_1' at 8 am with a positive latent heat have an average future lifetime in C_1' of 4 hours (Figure 16). Flows in C_1' at 6 pm with a positive latent heat have an average future lifetime in C_1' of 16 hours.

During the working hours active flows are large in number. They usually correspond to the traffic attracted by different users inside the network. During the night, traffic destined toward specific network prefixes is likely to be due to automated computer processes that transfer data across the network. This kind of processes is likely to be more persistent in time, since it is not affected by user demand, which is volatile. Therefore, flows detected as persistent heavy-hitters during the night could be expected to behave more persistently across time under such an assumption.

Very long-lived and very short-lived flows in C_1' are likely to skew the average value of a flow's future lifetime in C_1'. Therefore, we also look into the median value of the collected distribution in Figure 17. We notice that the median future lifetime of a flow with a latent heat greater than 60 Kbps at 8 am is 8 hours. This finding is a confirmation of the capability of our scheme to correctly detect persistent heavy-hitters, since these flows are likely to correspond to businesses attracting traffic over the net-

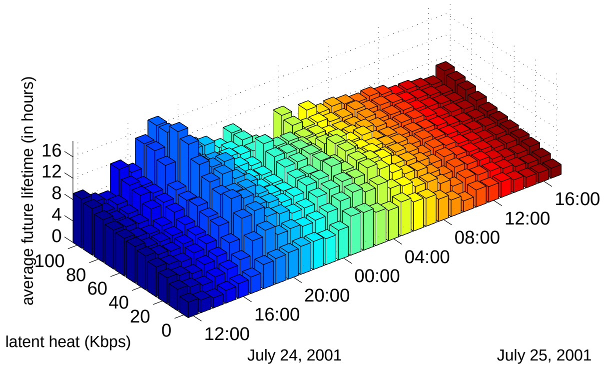

Fig. 16. Average future lifetime of a flow in C'_1 at different times throughout the day ($\gamma = 0.2$, $thres = 0$).

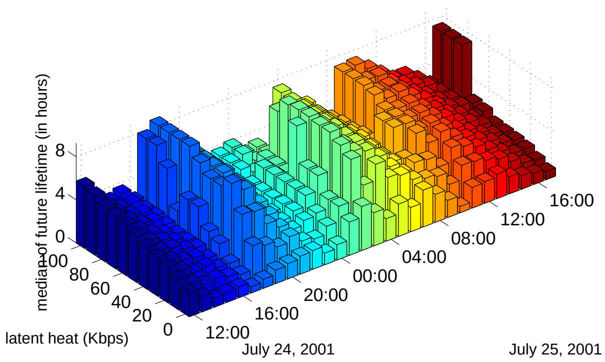

Fig. 17. Median future lifetime of a flow in C'_1 at different times throughout the day ($\gamma = 0.2$, $thres = 0$).

work during the 8 working hours of the day. The median future lifetime of a flow in C'_1 also peaks at 6 pm and 5 am. The former corresponds to the end of the working hours and the beginning of the night where flows have been shown to be more persistent. The latter corresponds to the beginning of the working hours on the east coast of the United States.

VIII. Conclusions

In this work we addressed the design principles of traffic engineering applications. For scalability reasons, traffic engineering applications typically select a small set of high-bandwidth flows that they treat differently inside the network in order to achieve a specific performance objective by affecting the majority of the traffic. Throughout the literature these flows are implicitly assumed to perform in a persistent fashion and thus are selected so that they capture the majority of the load at one point in time and treated as such for the timeframe at which the application operates [1], [8], [9], [11].

Based on data collected from an operational tier-1 IP backbone network we showed that flows defined at the level of a destination network prefix are very volatile in terms of volume. This volatility is manifested through short-lived changes in the flows' bandwidth and long-term changes that may reduce their bandwidth significantly or even lead to inactivity. As a direct result, the load captured by a specific set of flows at one point in time may significantly drop over time.

We showed that simple classification schemes that select heavy-hitter flows based on their behavior during a single time interval fail to capture the majority of the traffic through time. In addition the set of flows that they would select across time intervals would differ greatly from one interval to the next.

We proposed a novel classification scheme that incorporates historical flow behavior in the selection of heavy-hitter flows for traffic engineering purposes. We showed that our scheme is less sensitive to outliers, and leads to a more consistent set of flows selected across time intervals. More importantly, we demonstrated that heavy-hitter selection in the context of traffic engineering presents a tradeoff between the captured amount of traffic and the required frequency of the flow reselection across time. One cannot capture the majority of the load across time in the same set of flows. Thus, flows need to be reselected at frequent intervals to lead to the desirable end result. On the other hand, if the selected flows are chosen such that they account for a smaller amount of traffic (say one third of the overall load) then reselection can be performed less frequently (on the order of hours). Under both scenarios, our latent heat scheme was shown to outperform previously proposed schemes. Our concluding remark is that traffic engineering applications need to take flow volatility into account and be ready to sustain the overhead of frequent flow reselection if they intend to affect a large amount of traffic.

IX. Acknowledgment

We would like to thank Chen-Nee Chuah, Gianluca Iannaccone, Patrick Thiran and Supratik Bhattacharyya for their comments on previous versions of this paper. We would also like to thank Mark Crovella for providing us with the `aest` tool.

References

[1] H. Abrahamsson, B. Ahlgren, J. Alonso, A. Andersson, and P. Kreuger. A Multi Path Routing Algorithm for IP Networks Based on Flow Optimisation. In *International Workshop on Quality of future Internet Services (QofIS)*, Zurich, Switzerland, October 2002.

[2] M. Crovella and M. Taqqu. Estimating the Heavy Tail Index from Scaling Properties. *Methodology and Computing in Applied Probability*, 1999.

[3] C. Estan and G. Varghese. New Directions in Traffic Measurement and Accounting. In *ACM SIGCOMM*, Pittsburgh, PA, USA, August 2002.

[4] W. Fang and L. Peterson. Inter-AS Traffic Patterns and Their Implications. *IEEE Globecom*, December 1999. Brazil.

[5] C. Fraleigh, C. Diot, B. Lyles, S. Moon, P. Owezarski, K. Papagiannaki, and F. Tobagi. Design and Deployment of a Passive Monitoring Infrastructure. In *Passive and Active Measurement Workshop*, Amsterdam, April 2001.

[6] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang. BGP Routing Stability of Popular Destinations. In *ACM SIGCOMM Internet Measurement Workshop*, Marseilles, France, November 2002.

[7] A. Shaikh, J. Rexford, and K. Shin. Load-Sensitive Routing of Long-Lived IP Flows. *ACM SIGCOMM*, September 1999.

[8] A. Sridharan, R. Guérin, and C. Diot. Achieving Near-Optimal Traffic Engineering Solutions for Current OSPF/IS-IS Networks. In *IEEE Infocom*, San Francisco, March 2003.

[9] X. Su and G. de Veciana. Dynamic multi-path routing: asymptotic approximation and simulations. *ACM SIGMETRICS Performance Evaluation Review*, 29(1):25–36, June 2001.

[10] G. Trotter. Terminology for Forwarding Information Base (FIB) based Router Performance. RFC 3222, December 2001.

[11] S. Uhlig and O. Bonaventure. On the Cost of Using MPLS for Interdomain Traffic. *Quality of Future Internet Services*, September 2000. Berlin.
[12] Y. Zhang, L. Breslau, V. Paxson, and S. Shenker. On the Characteristics and Origins of Internet Flow Rates. In *ACM SIGCOMM*, Pittsburgh, PA, USA, August 2002.

APPENDIX

In time series analysis and forecasting, we use a *Moving Average* model to estimate the current value of the mean based on past measurements. The moving average is often called a "smoothed" version of the original series, since short-term averaging has the effect of smoothing out the bumps in the original series.

At any instant k, the average of the latest n samples of a data sequence x_i is given by

$$x_k = \frac{1}{n} \sum_{i=k-n}^{k-1} x_i.$$

Hence, the forecast equals the simple average of the last n observations. This average is "centered" at period $i = (n+1)/2$ which implies that the estimate of the local mean will tend to lag behind the true value of the local mean by about $(n+1)/2$ periods. Thus, we say that the *average age* of the data in the simple moving average is $(n+1)/2$ relative to the period for which the metric is computed.

The simple moving average model has the undesirable property that it treats the last n observations equally and completely ignores all preceding observations. Intuitively, past data should be discounted in a more gradual fashion. The *simple exponential smoothing (SES)* model accomplishes this. If α denotes the "smoothing constant", with $0 < \alpha \leq 1$, and s_t denotes the value of the smoothed series at period t, the following formula can be used recursively to update the smoothed series as new observations are recorded:

$$s_t = \alpha x_{t-1} + (1-\alpha) s_{t-1},\; 0 < \alpha \leq 1,\; and\, t \geq 3 \qquad (5)$$

Thus, the current smoothed value is an interpolation between the previous smoothed value and the previous observation, where α controls the closeness of the interpolated value to the most recent observation. If we expand Equation (5) by substituting for s_{i-1} we have:

$$\begin{aligned} s_t &= \alpha x_{t-1} + (1-\alpha)\left[\alpha x_{t-2} + (1-\alpha) s_{t-2}\right] \\ &= \alpha x_{t-1} + \alpha(1-\alpha) x_{t-2} + (1-\alpha)^2 s_{t-2} \end{aligned}$$

By recursively substituting for s_{t-2}, then for s_{t-3} and so forth until we reach s_2 which is equal to x_1, it can be shown that the expanding equation can be written as:

$$s_t = \alpha \sum_{i=1}^{t-2} (1-\alpha)^{i-1} x_{t-i} + (1-\alpha)^{t-2} s_2,\; t \geq 3\; and\, 0 < \alpha \leq 1$$

The above equation shows the exponential behaviour. We further can see from the summation term that the contribution of each value $x_i,\; i < t$ to the smoothed value s_t becomes less at each consecutive time interval. The speed at which old responses are dampened is a function of the value α. When α is close to 1, dampening is quick and when α is close to 0, dampening is slow.

The weights $\alpha(1-\alpha)^t$ decrease geometrically. Given that the mean of the geometric distribution of the weights is equal to $1/\alpha$, the "average age" of the data in the simple exponential smoothing forecast is $1/\alpha$ relative to the period for which the forecast is computed. We showed that the "average age" of the data in a simple moving average estimate for a period of n intervals is $(n+1)/2$. Consequently, the rule of thumb for the setting of the "smoothing constant" α in order to exhibit an average age of n periods can be described as:

$$\frac{1}{\alpha} \approx \frac{n+1}{2},\; or\; \alpha \approx \frac{2}{n+1} \qquad (6)$$