Designing Traffic Profiles for Bursty Internet Traffic

Xiaowei Yang (yxw@lcs.mit.edu)
MIT Laboratory for Computer Science
Cambridge, MA 02139

Abstract—**This paper proposes a new class of traffic profiles that is better suited for metering bursty Internet traffic streams than the traditional token bucket profile. A good traffic profile should satisfy two criteria: first, it should consider packets from a conforming traffic stream as in-profile with high probability to ensure a strong QoS guarantee; second, it should limit the network resources consumed by a non-conforming traffic stream to no more than that consumed by a conforming stream. We model a bursty Internet traffic stream as an ON/OFF stream, where both the ON-period and the OFF-period have a heavy-tailed distribution. Our study shows that the heavy-tailed distribution leads to an excessive randomness in the long-term session rate distribution. Therefore, it is inherently difficult for any profile that limits the long-term average session rate to give a strong QoS guarantee for the conforming traffic streams. Our simulation demonstrates that a token bucket profile that couples the average rate control and the burst size control has a weak QoS guarantee. Based on this result, we propose a new class of traffic profiles that decouples the long term average rate control from the burst size control. Compared to a token bucket profile, this profile improves the level of QoS for a conforming traffic stream, yet limits the "effective bandwidth" consumed by a non-conforming traffic stream.**

I. Introduction

This paper discusses the design of traffic profiles for bursty Internet traffic. We propose a new class of traffic profiles that performs better than the traditional token bucket profile. In the Diffserv [1] architecture, a traffic profile defines the rules for a meter to decide whether a packet from a traffic stream is in-profile or out-of-profile. The packet may be subject to different conditioning actions. When there is a Service Level Agreement (SLA) present, the essential goal of a profile is not to smooth bursty traffic. Instead, it is to check whether a traffic stream conforms to the SLA. A good traffic profile needs to satisfy two criteria. First, from the perspective of providing a QoS guarantee, it is desirable to have a high probability guarantee that packets from a conforming stream will be identified as in-profile. A conforming stream is a traffic stream whose statistics are allowed by a traffic profile. Second, from the perspective of policing against misbehavior, a traffic profile should limit the "effective bandwidth" [2] consumed by the non-conforming stream to be no more than that consumed by a conforming stream. Thus the misbehavior of the non-conforming stream will not degrade the level of service for a conforming stream. An example of a non-conforming stream is one that has a much higher expected transmission rate than that allowed by a traffic profile.

It is desirable to have a range of traffic profiles for different classes of traffic, and for differentiation within a class. For example, traffic streams generated by constant bit rate video applications have quite different statistics than those generated by web browsing sessions, thus require different profiles to meet their QoS requirements. It is worth noting that the type of traffic we consider in this paper is most often generated by document transfers, such as web browsing. Traditionally, the QoS requirement for such traffic (a.k.a. elastic traffic) is less obvious than that for the real-time traffic. In this paper, the QoS requirement we consider is *transparency* [3]. If a burst is sent at the specified peak rate, we say the transfer is *transparent*. If a burst is not sent at the specified peak rate, a user will notice delay, spend time waiting, and become dissatisfied. Thus, we consider *transparency* as a valid QoS requirement for such interactive applications. We admit that the ultimate QoS measurement is user satisfaction, which can not be completely determined by the measurement of network performance. For an interactive traffic stream, a user may value some bursts more than others. Since we do not have a precise method to quantify user satisfaction, as an approximation, we assume that the higher the percentage of bursts that are sent transparently , and the more the bytes that are sent transparently, the higher the user's satisfaction is. There have been related studies on traffic profiling algorithms, such as different window policing mechanisms [4], multiple token banks [5], and the *fractal leaky bucket* [6] algorithm.

In this paper, we model a bursty Internet traffic stream as an ON/OFF stream, where both the ON-period and the OFF-period have heavy-tailed distributions. This model is a simplified derivation from recent results of traffic measurement and modeling [7], [8], [9]. According to Barford and Crovella's work [10], users' Web browsing behaviors follow the ON/OFF pattern, where ON-period has a heavy-tailed distribution attributed to the heavy-tailed file size distribution.

Our simulation and analysis show that there is no "natural" average session rate for a single traffic stream. Thus, it is unlikely for any profile that limits the long-term average session rate to give a strong QoS guarantee. Our simulation also demonstrates that a token bucket profile that couples the average rate control and the burst size control has a weak QoS guarantee. Based on this result, we propose a new class of traffic profiles that decouples the long term average rate control from the burst size control. Compared to a token bucket profile, this profile improves the level of QoS for a conforming stream, yet effectively limits the "effective bandwidth" consumed by a non-conforming traffic stream.

II. Traffic Model, Simulation Method and Notation

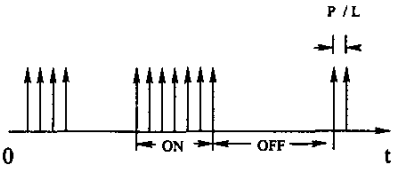

Fig. 1. The ON/OFF Traffic Model.

In this section, we describe the traffic model, the simulation method and some notation. In our simulation, a bursty traffic stream from a

This work was sponsored by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-00-2-0553.

0-7803-7632-3/02/$17.00 ©2002 IEEE

TABLE I
SIMULATION PARAMETERS

Period	*Shape* (α)	*Location* (β)	*Distribution*	*Mean*
ON	1.05	1024 Bytes	$P\{X < x\} = 1 - (\frac{1024}{x})^{1.05}, x > 1024$	14026 Bytes
OFF	1.4	1 Second	$P\{X < x\} = 1 - (\frac{1}{x})^{1.4}, x > 1$	3.4939 Seconds
Packet Length		1024 Bytes		
Peak Rate		1.5Mbps		

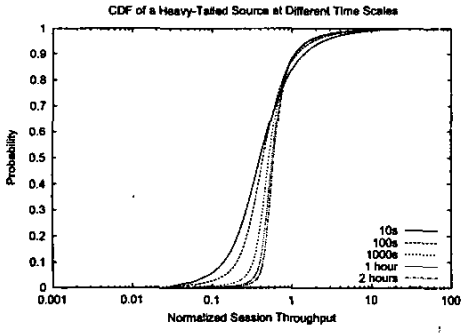

Fig. 2. The convergence of a Pareto ON/OFF Traffic Stream. The different curves are for session lengths of 10s, 100s, 1000s, 1 hour and 2 hours.

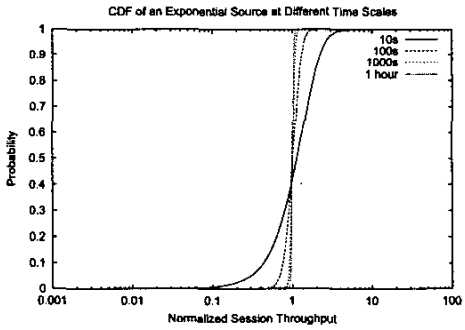

Fig. 3. The convergence of an Exponential ON/OFF Traffic Stream. The different curves are for session lengths of 10s, 100s, 1000s, 1 hour and 2 hours.

single source is modeled as an ON/OFF traffic stream. The ON-period models a single flow, such as the transfer of a single web page, and the OFF-period models users' thinking time. ON-periods and OFF-periods are strictly alternating. The stream models traffic generated during a session, i.e., a busy period of user activities, such as web browsing for an hour. For bursty Internet traffic, both ON-periods and OFF-periods are modeled by Pareto distributions. In an ON-period, packet arrivals are spaced out by P/L, where P is the packet size and L is the peak rate. Figure 1 depicts the traffic model. In our simulation, the traffic stream is passed through a meter that decides whether a packet is in-profile or out-of-profile according to a traffic profile. As we are interested in the *transparency* QoS requirement, out-of-profile traffic is shaped, instead of dropped or marked. Thus, the transfer time of a burst can be used to measure performance.

Traditional traffic models assume that the burst size is exponentially distributed. To better understand the difference between these two models, we simulate both the Exponential distribution and the Pareto distribution in some experiments, and compare the results. If a Pareto distribution has mean μ and location parameter m, the corresponding exponential distribution is defined as $P\{X < x\} = 1 - e^{\frac{x-m}{\mu-m}}, x > m$.

In this paper, a token bucket profile with a token rate r bytes/s and a token depth b bytes is denoted by the pair (r, b). We always assume our simulation starts at time 0. And $A(t)$ denotes the cumulative arrivals of a traffic stream up to time t. Simulation results shown are obtained using the parameters in Table I unless otherwise specified. These parameters for the Pareto distribution are chosen based on recent research results on traffic measurement and modeling [10], [9].

III. AVERAGE SESSION RATE OF A BURSTY SOURCE

For a heavy-tailed ON-period distribution, there is no "natural" burst size [7], to which the token bucket size can be set. In this section, we look at the distribution of the average session rate of a traffic stream at different time scales via simulation. A session is simulated by a single run with some length T. Each distribution curve is sampled from 10^6 runs. Figure 2 and Figure 3 show the cumulative distribution function of the normalized session rate of a Pareto distribution and an Exponential distribution over multiple time scales. A normalized session rate is computed by dividing the average rate in a simulated session by the expected rate of the ON/OFF stream. The tail is shortened quickly for the Exponential stream, which shows that the throughput of an exponential stream averaging over an hour duration has a small variance and there exists a "natural" session rate to which the token rate can be set. In contrast, the tail shape for the Pareto distribution does not change much for simulations ranging from ten seconds to two hours, which suggests that even averaging over a duration as long as two hours, the session rate of a heavy-tailed ON/OFF stream has a large variance. Hence it has non-negligible probabilities of reaching quite large or quite small values. The observed maximum throughput is only limited by the peak rate. Intuitively, the simulation results can be explained by the central limit theorem and the corresponding limit theorem for random variables with infinite variance [11], [12]. For a detailed analysis, see [13].

This slow-convergence property poses a problem for designing traffic profiles. A bursty traffic stream generated by a particular application is a random process. We hope to predict its behavior from prior observations in order to negotiate a proper profile for future usage. For an Exponential ON/OFF traffic stream, we observe that it is possible to predict the stream's long term behavior. However, for a heavy-tailed ON/OFF stream, even if we assume that the application always generates traffic conforming to a known distribution, the behavior of an individual session is still not predictable. For any profile that limits the average session rate to r, the number of arrived tokens $A(t)$ in a duration t is limited by rt. Both our simulation and analysis demon-

strate that for any time scale t, there is a non-negligible probability that an application may generate near-peak traffic for the entire duration t. Hence, unless $A(t)$ is close to Lt (recall that L is the peak rate), an application will generate out-of-profile traffic with non-negligible probability, and we can not guarantee that its QoS requirements will be satisfied.

IV. Evaluating the Performance of Token Bucket Profiles

In this section, we look at the performance of a token bucket profile. A token bucket profile (r, b) limits the maximum burst size of a traffic stream to b, as well as the average rate of a session with length t to $r + b/t$, thus enhances the predicability of the output traffic stream. It has no long term memory. When the bucket is full, extra tokens are discarded, and the good behavior of a stream is forgotten. The decision made on a packet depends only on the short term history of a traffic stream, i.e., whether the previous burst has consumed all the tokens, but not on the long term history.

A. Performance of a Token Bucket Profile

We simulate the case where an ON/OFF traffic stream is policed against a token bucket profile. When the transfer of a burst is finished, an OFF time follows before a new burst arrives. The maximum bucket size is set at ten times the mean burst size. At the beginning of each simulation, the bucket is full. For each run, we record the time it takes for each burst to be transmitted. The time is normalized using the peak-rate transfer time of each burst. We then draw a distribution of the normalized transfer time, both in the number of bursts and in the number of bytes. Each distribution is sampled from 10,000 runs. Each simulation runs for the duration which it takes to transmit the amount of traffic that would be sent in a one-hour session with no traffic profile present. Figure 4 and Figure 5 show the performance of a token bucket profile for shaping a Pareto ON/OFF source and an Exponential ON/OFF source. The token rate varies from the expected session rate to twice the expected session rate.

As seen from the figures, when the token rate is set to the expected session rate, for a Pareto ON/OFF stream, more than 98% of bursts can be sent without distortion. However, those bursts only consist of 45% of the total bytes sent. For an Exponential ON/OFF stream, more than 90% of files and 80% of bytes can be sent without distortion. When the token rate is increased to twice the average session rate, the performance for the Exponential stream is improved much more significantly than that for the Pareto stream. This result can be predicted from the session rate distribution, as a quite high session rate may happen for a Pareto ON/OFF stream with non-negligible probability. The session rate may be very close to the peak rate. Hence increasing the token rate does not increase the performance significantly for the Pareto ON/OFF stream.

B. How to Improve QoS Guarantees

A token bucket profile limits both the burst size and the average session rate. We can either increase the limit on the average session rate or increase the limit on the maximum burst size to improve QoS guarantees. To formulate ideas on how to improve the performance of a token bucket profile for a heavy-tailed ON/OFF stream, we compare the effectiveness of increasing the bucket size and increasing the token rate from three aspects. First, we compare how effective they are in improving the QoS guarantees for a single conforming traffic stream. By "conforming", we mean that the traffic stream is generated from the pre-negotiated distribution. Second, we compare how effective they are in limiting the QoS guarantees for a single aggressive stream. By "aggressive", we mean that the traffic stream consumes more bandwidth than pre-negotiated. Thirdly, we compare how effective they are in controlling aggregated aggressive streams. An effective profile should limit the peak bandwidth of the multiplexed aggressive streams to be no more than the peak bandwidth of the multiplexed conforming streams without traffic profiles. Thus, in the worse case, when the majority of the users are misbehaving, a network provider can still provide QoS for the in-profile traffic, given that he has provisioned his network according to the peak bandwidth of the multiplexed conforming streams.

To conduct the first comparison, we run simulations with a token bucket profile that allows unbounded bucket size. And the initial number of tokens in the bucket is set at ten times the average burst size. The result from this profile is compared against the normal token bucket profile (r, b) and $(2r, b)$, respectively. Figure 6 and Figure 7 show the performance comparison when the stream's expected rate is r. The token bucket profile $(2r, b)$ shortens the token-rate burst transfer time by half, but it does not significantly increase the number of bursts or bytes that are sent at the peak rate. In contrast, the profile $(r, unlimited)$ reduces the non-peak rate transfered files from 1.7% to less than 0.2% and the non-peak rate transfered bytes from 55% to 33%.

For the second comparison, we set the minimum burst size of an aggressive stream at ten times that of a conforming stream. Such an aggressive stream is policed by profiles (r, b), $(2r, b)$ and $(r, unlimited)$, where r is the expected rate of a conforming stream and b is ten times the mean burst size of a conforming stream. The results are shown in Figure 8 and Figure 9. As can be seen, the profile $(2r, b)$ is not as effective as the profile $(r, unlimited)$. The profile $(2r, b)$ allows more bursts and bytes to be sent without distortion or with less distortion and can not limit the long term session rate to be r. Thus, it is not as effective as the profile $(r, unlimited)$ at discriminating conforming and aggressive streams.

For the third comparison, we compare the normalized peak data rate of aggregated aggressive traffic streams when shaped by the three profiles (r, b), $(2r, b)$ and $(r, unlimited)$. A normalized peak data rate is computed as P/nr, where P is the measured average peak rate of aggregated streams, r is the expected data rate of a single conforming traffic stream, and n is the number of streams. The aggregated peak rate P_r is a random variable, with an upper bound $n \times L$, where L is the peak data rate of a single traffic stream. In our simulation, all streams are homogeneous. For each n, we ran $10000/n$ simulations. Each of them simulates an 8-hour trace. We record the peak arrival rate seen in $5ms$ intervals in the trace of each run. The length of the interval is chosen because the spacing between packets from a single source is about $5ms$ in our simulation. We then compute the average peak rate P and the standard deviation of P over $10000/n$ runs. As we do not know the distribution of P_r and each simulation is run for a long time, we use P as a representative value such that with a high probability, the peak rate of the aggregated traffic streams P_r is less than P.

Figure 10 shows the simulation results, together with the normalized data rate for conforming traffic streams when there is no profile and no shaping. As shown in Figure 10, the normalized data rate degenerates as the number of multiplexed sources increases, indicating

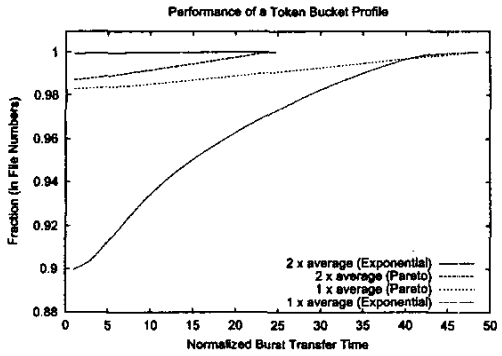

Fig. 4. Fraction counted by the number of bursts for an ON/OFF traffic stream. When the token rate doubles, most bursts of the Exponential stream are sent at the peak rate.

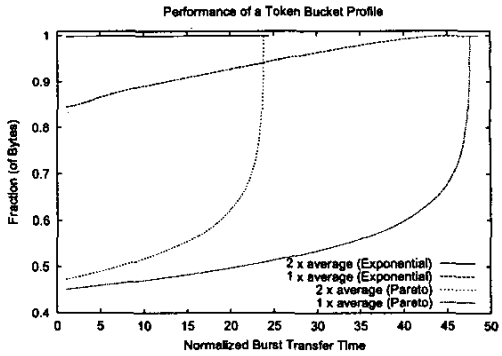

Fig. 5. Fraction counted by bytes. When the token rate doubles, most bytes of the Exponential stream are sent at the peak rate. For the heavy-tailed traffic stream, even more than 98% of the bursts are sent at the peak rate, they constitute merely about 45% of the bytes.

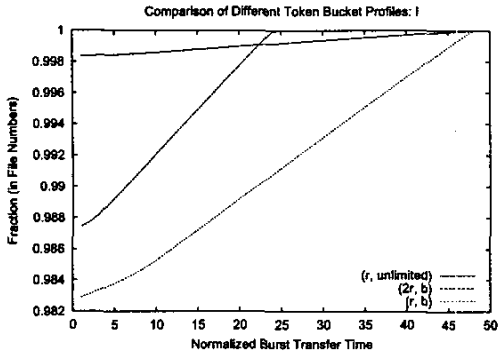

Fig. 6. Fraction counted by the number of bursts.

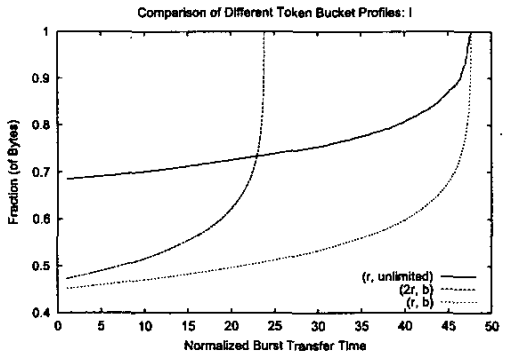

Fig. 7. Fraction counted by bytes.

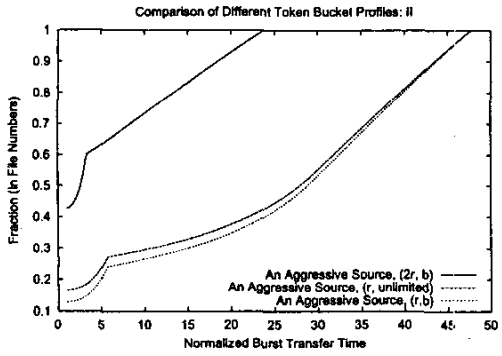

Fig. 8. Fraction counted by the number of bursts.

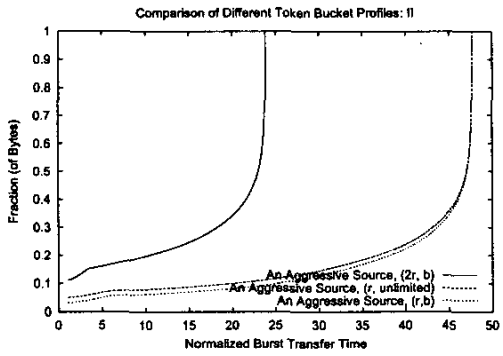

Fig. 9. Fraction counted by bytes.

a statistical multiplexing gain for bursty traffic. The normalized peak data rates of the aggregated aggressive streams shaped by the profile $(r, unlimited)$ and (r, b) are quite close to that of the aggregated conforming streams without shaping. For the profile $(2r, b)$, the aggregated normalized peak data rate exceeds that of conforming streams, which means that the profile is not effective in limiting the effective bandwidth of aggressive streams. A network will become congested when the peak of the aggregated traffic exceeds the bandwidth provisioned according to the conforming sources without profiles, and thus can not guarantee the QoS for the in-profile traffic.

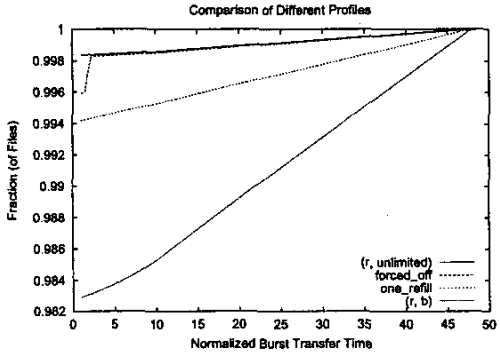

Fig. 11. Fraction counted by the number of bursts.

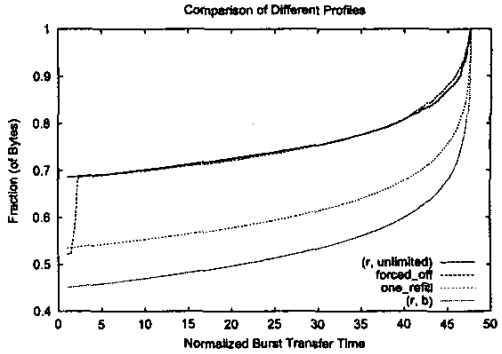

Fig. 12. Fraction counted by bytes.

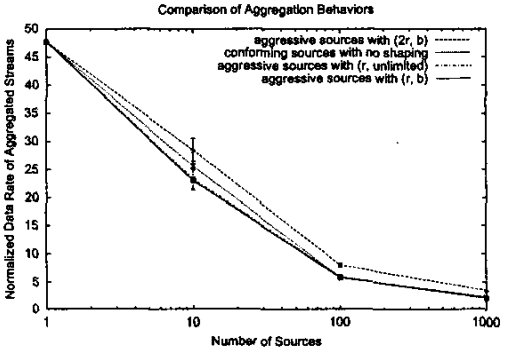

Fig. 10. The normalized data rate of aggregated traffic when applying no profile to conforming traffic streams, profile (r, b), $(2r, b)$ and $(r, unlimited)$ to aggressive traffic streams.

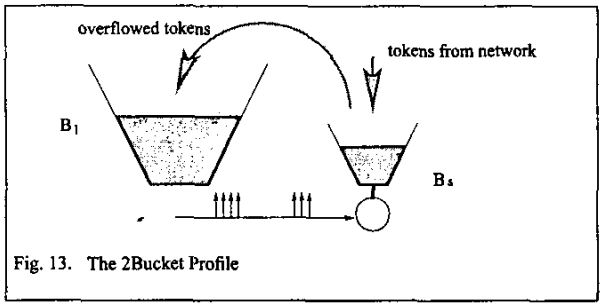

Fig. 13. The 2Bucket Profile

V. A New Traffic Profile: the 2Bucket Profile

From the above experiments, we conclude that to effectively improve QoS without compromising discrimination against misbehavior, it is better to increase the token bucket size than to increase the token rate. However, the $(r, unlimited)$ profile offers no control over burst size. A malicious user may keep silent for almost the entire course of a session and then burst out all his accumulated tokens. This observation leads to our design of a new class of traffic profile: the 2Bucket profile. A visualization of the profile is shown in Figure 13.

One bucket B_s that has a maximum capacity b is used to control the burst size. The other bucket B_l, which has an unlimited capacity, is used to record the long term history. Tokens arriving from the network are stored into B_s first. If B_s is overflowed, tokens are dumped into B_l. When a packet arrives, if B_s has enough tokens, the packet is transmitted immediately. If B_s is empty and B_l has extra tokens, a refill process may be triggered. This process is triggered according to a refill policy. The choice of refill policies is a key design parameter of the 2Bucket profile. In a refill, tokens are taken from B_l to fill B_s. After B_s is filled with enough tokens, the packet may be sent. If both B_s and B_l are empty, or B_s is empty and a new refill process is not allowed, the packet has to wait for tokens to come from the network. A refill policy decides when and how tokens are refilled from B_l into B_s. We use the terms B_s and B_l to refer to both the available tokens in the two buckets and the buckets themselves, as the meanings will be clear from context. Recall that $A(t)$ is the cumulative arrivals from a stream in a time period t. This two-bucket profile ensures that as long as $A(t) < rt + b$, $B_l + B_s > 0$.

The behavior of the 2Bucket profile is quite different when different refill policies are deployed. A refill policy determines the burstiness allowed by the profile. We will describe two possible policies and analyze their performance. It is also possible to design other refill policies, depending on the type of controls we need. The first policy is called "*forced_off*". For this policy, a refill process fills B_s with $\min(b, B_l)$ tokens at the peak rate. Two consecutive refills must be separated by at least the time interval *forced_off*. This policy will break a long burst into several small bursts, as long as the long term behavior of a traffic stream is good, i.e., $B_l + B_s$ is not empty. If the long burst is a document transfer, this break-down will increase the transfer time. Depending on the values of b and *forced_off*, the transfer time can still be much less than the token-rate transfer time. If the burst is actually a video stream (for example, a user asks for a cheap bursty-traffic profile but uses it to watch on-line video), this break-down will give intolerable performance. Thus, it can discourage the misuse of a profile.

We call the second policy *one_refill*. For this policy, refill is more strict. Only one refill of $\min(b, B_l)$ tokens is allowed in any ON period. It limits the maximum burst size to be at most $2 \times b$. If a stream's long term behavior is good, each ON-period is guaranteed a refill. In contrast, for a token-bucket profile, if the previous burst exhausts all available tokens in the bucket, the next burst has to wait for more to-

kens.

Performance of this new profile with two policies is shown in Figure 11 and Figure 12, along with the corresponding token-bucket profiles (r, b) and $(r, unlimited)$. In our simulation, for the *forced_off* policy, the parameter *forced_off* is set to the location parameter of OFF-period distribution, *min_off*. For the *one_refill* policy, if we have not seen a packet from a source for a time period *min_off*, we conclude one burst has ended and the stream is eligible for a new refill. The profile $(r, unlimited)$ is actually a special case of the 2Bucket profile. When the parameter *forced_off* is set to zero, the 2Bucket profile behaves exactly the same as $(r, unlimited)$.

Notice that there is a sharp knee in the curve for *forced_off*, and the performance of *forced_off* and $(r, unlimited)$ are almost the same after the knee. When examining the raw transfer time recorded in the trace file, we found that the bursts after the knee were sent when B_l ran out of tokens. Therefore, the rest of the bursts can only be sent at the token rate, despite the refill policies. When B_l runs out of tokens, the long term behavior of a stream reaches the limitation imposed by the token rate. Clearly neither *forced_off* nor $(r, unlimited)$ can improve the performance in this case. The steep region of the curve consists of bursts that are transmitted when B_l has extra tokens. In the $(r, unlimited)$ case, bursts are sent at the peak rate immediately. In the case of *forced_off*, a *forced_off* token-rate period is forced between two bursts. The transfer time is slightly lengthened under this policy. Compared to the token bucket profile, both *forced_off* and *one_refill* improve performance.

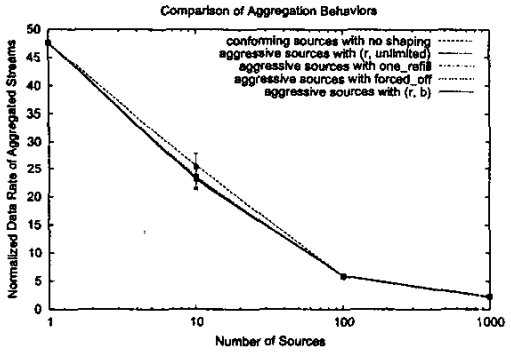

Fig. 14. The normalized data rate of aggregated traffic with profile $(r, unlimited)$, *one_refill*, *forced_off* and (r, b) applied to aggressive traffic streams, and the normalized peak data rate for conforming traffic streams when there is no profile and no shaping. The performance of each profile is similar.

Figure 14 shows the normalized peak data rate of aggregated aggressive traffic streams when shaped by the four profiles (r, b), $(r, unlimited)$, *one_refill* and *forced_off*, and the normalized peak data rate for conforming traffic streams when there is no profile and no shaping. The four profiles perform indistinguishably. All of them are able to limit the normalized peak data rate of the aggregated aggregated streams quite close to that of the aggregated conforming streams. For details, see [13].

VI. Conclusions and Future Work

In this paper, we describe a new class of traffic profiles — the 2Bucket profile. When metering bursty traffic, this profile achieves better performance than the traditional token bucket profile by decoupling the long term average rate control from the burst size control. Not only is the profile able to remember the long term behavior of a traffic stream, it also enables the network to control the burst size. We also present two possible refill policies that can be applied to the 2Bucket profile to demonstrate how a network can use the profile to enforce different QoS policies.

In the future, we intend to build a prototype system and test the effectiveness of the 2Bucket profile through user experiments and trace-driven simulation. In particular, we are interested in testing how different policies and parameter settings affect users' QoS assessments.

VII. Acknowledgment

I am grateful to my advisor David Clark for his guidance and support. I thank Arthur Berger and John Wroclawski for helpful discussions and comments on early drafts of the paper. I also thank the anonymous reviewers for their perceptive comments.

References

[1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An Architecture for Differentiated Services," December 1998, RFC 2475.

[2] Frank Kelly, "Notes on Effective Bandwidths," in *Stochastic Networks*, Frank Kelly, S. Zachary, and I. Ziedins, Eds., pp. 141–168. Oxford Science Publications, 1996.

[3] J.W. Roberts, "Engineering for Quality of Service," in *Self-Similar Network Traffic and Performance Evaluation*, Kihong Park and Walter Willinger, Eds., pp. 401–420. John Wiley & Sons, Inc., 2000.

[4] Erwin Rathgeb, "Modeling and Performance Comparison of Policing Mechanisms for ATM Networks," *IEEE Journal on Selected Areas in Communications*, vol. 9, no. 3, pp. 325–334, April 1991.

[5] Arthur W. Berger and Ward Whitt, "A Multi-Class Input-Regulation Throttle," in *Proceedings of the 29th Conference on Decision and Control*, Honolulu, Hawaii, December 1990.

[6] Nelson L. S. Fonseca, Gilberto S. Mayor, and Cesar A.V. Neto, "On the Equivalent Bandwidth of Self-Similar Sources," *ACM Transactions on Modeling and Computer Simulation*, vol. 10, no. 2, pp. 104–124, 2000.

[7] Vern Paxson and Sally Floyd, "Wide area traffic: the failure of Poisson modeling," *IEEE/ACM Transactions on Networking*, vol. 3, no. 3, pp. 226–244, 1995.

[8] Mark E. Crovella and Azer Bestavros, "Self-similarity in World Wide Web traffic: evidence and possible causes," *IEEE/ACM Transactions on Networking*, vol. 5, no. 6, pp. 835–846, 1997.

[9] Will E. Leland, Murad S. Taqqu, Walter Willinger, and Daniel V. Wilson, "On the Self-Similar Nature of Ethernet Traffic," *IEEE/ACM Transactions on Networking*, vol. 2, no. 1, pp. 354–367, February 1994.

[10] Paul Barford and Mark Crovella, "Generating Representative Web Workloads for Network and Server Performance Evaluation," in *Proceedings of ACM SIGMETRICS'97*, June 1998, pp. 151–160.

[11] William Feller, *An Introduction to Probability Theory and Its Applications*, vol. II, chapter VI, John Wiley & Sons, 1971.

[12] Mark Crovella and Lester Lipsky, "Long-lasting transient conditions in simulations with heavy-tailed workloads," in *Winter Simulation Conference*, 1997, pp. 1005–1012.

[13] Xiaowei Yang, "Designing Traffic Profiles for Bursty Internet Traffic," Tech. Rep., MIT/LCS/TR-860, August 2001.