

MPAT: Aggregate TCP Congestion Management as a Building Block for Internet QoS

Manpreet Singh, Prashant Pradhan* and Paul Francis
Cornell University, *IBM T.J. Watson Research Center
{manpreet,francis}@cs.cornell.edu, ppradhan@us.ibm.com

Abstract

Today Internet QoS is deployed piecemeal—typically at known bottleneck links like enterprise access links or wireless links. A more comprehensive, end-to-end QoS deployment, for instance across large enterprise networks or the global Internet, remains elusive. There is growing interest in the idea of using overlay networks to provide differential QoS services (improve performance for some flows at the expense of other flows). A necessary building block is the ability to provide differential service over a single overlay link that traverses many IP router hops. This paper presents MPAT, the first truly scalable algorithm for fairly providing differential services to TCP flows that share a bottleneck link. Unlike known schemes, our approach preserves the cumulative fair share of the aggregated flows even where the number of flows in the aggregate is large. Specifically we demonstrate, primarily through experiments on the real Internet, that congestion state can be shared across more than 100 TCP flows with throughput differentials of 95:1. This is up to five times better than differentials achievable by known techniques. Indeed, MPAT scalability is limited only by the delay-bandwidth product of the aggregated flows. With this tool, it is now possible to seriously explore the viability of network QoS through overlay network services.

1 Introduction

After the heady optimism of the vision of a universal QoS in the form of RSVP [29, 16], we have come to understand that realistically QoS can only be deployed piecemeal. Rather than a single ubiquitous QoS architecture, what is evolving is an arsenal of tools that can be opportunistically deployed where benefits can clearly be demonstrated. These benefits are most clear where a bottleneck resource (e.g. a single bottleneck link) can be identified, there is an identifiable set of user or traffic classes that would benefit from arbitration of that bottleneck, and the organization with a vested interest in providing

QoS has control over the bottleneck resource.

One example is an enterprise network using IP telephony over the Internet[34], where the bottleneck is the access link between the enterprise locations and the Internet. Differentiation between voice and non-voice traffic provides clear benefits, and this traffic can be identified by a router, say by looking for RTP packets. The enterprise has control over the routers on either end of the access link, either directly or by coordinating with its ISP. Another common example is in wireless networking, either 802.11 or cellular, where there may be different types of traffic or different classes of users (gold, silver, bronze customers, or emergency services) vying for scarce radio spectrum. Here, QoS mechanisms may be implemented at several layers in the the protocol stack (for instance, [14]), and the wireless network provider controls these protocols.

In these diverse examples, a single organization (enterprise, wireless network provider) was able to apply a QoS tool (router queuing, radio access control) to an identifiable bottleneck resource (enterprise access link, radio spectrum). These examples represent the “low hanging fruit”, if you will, of QoS tools—those cases where the opportunity and motivation are clear. What about the more difficult scenarios where there is an underlying network of some complexity, and no single organization has the motivation or wherewithal to implement QoS services in that network?

An interesting approach to create a network with controllable properties, on top of an unstructured Internet, is the overlay network. Overlay networks have been produced for reducing network delays [22], improving network resilience [6], and providing QoS to network applications like games [9]. The basic building block for providing QoS over an overlay network is the ability to provide QoS over each logical overlay link. Flows between the two ends of an overlay link share the same underlying network path, thus sharing the bottleneck link. Hence mechanisms to arbitrate the bottleneck link bandwidth to provide QoS can be deployed at the ends of the overlay link.

Note that the question of whether overlay networks is a valid

approach to providing QoS is an open question, and this paper does not answer that question. In particular, if the overlay changes the characteristics of the individual TCP flows that pass through it, and if multiple overlays try to “compete” under these circumstances, then the overlays may lose their effectiveness, or worse may destabilize the Internet. What is clear, however, is that if we are to make progress in understanding what can legitimately be achieved with overlays, we need to understand the characteristics of a single hop in the overlay. More specifically, in order to co-exist with Internet traffic (and with other similar overlays), such a QoS mechanism should satisfy certain properties. We state these properties below, assuming there are N flows belonging to an overlay using a logical overlay link.

- *Fairness*: The overlay flows should not steal bandwidth away from other flows (overlay or non-overlay) sharing the same bottleneck. In particular, other TCP or TCP-friendly flows using that bottleneck link should get the same bandwidth as they would with N standard TCP flows using the overlay. This requirement ensures that the overlay’s QoS mechanism is “transparent” to the background flows.
- *Utilization*: The overlay flows should be able to hold the fair share of the bottleneck link bandwidth entitled to them. In particular, the total bandwidth available to the overlay’s flows should be equal to the total fair share of N TCP flows. This ensures that the overlay flows do not lose bandwidth to the background traffic.
- *Scalability*: The above properties should hold for large values of N .

The key technical contribution of this paper is an aggregate TCP congestion management scheme that scales well by the number of flows in the aggregate (we have demonstrated up to $N=100$), that provides large differentiation ratios (we have demonstrated up to 95:1), and that does so fairly. We call the scheme MPAT, for Multi-Probe Aggregate TCP, because it maintains an active AIMD loop (i.e. a “probe”) for every TCP flow in the aggregate. This can be contrasted with multTCP[3], which tries to obtain N fair shares by running a single, more aggressive AIMD loop. A secondary contribution of this paper is to provide more experimental data and insights on the existing schemes (primarily multTCP). Even though the scalability and stability of multTCP is known to be limited, we feel it is important to provide this data for two reasons. First, contrasting MPAT with multTCP allows us to clarify and verify our intuitions as to why MPAT performs well. Second, multTCP, in spite of its stated limitations, is nevertheless still being proposed for overlay QoS [9] and so can be considered the incum-

bant. Thus it is important to make more direct comparisons between MPAT and multTCP.

The rest of this paper is organized as follows. Section 2 discusses the existing techniques to provide QoS to TCP or TFRC flows over a single path or bottleneck link, and demonstrates that they fall short of meeting the desirable criteria for such a QoS mechanism. Section 3 presents the proposed aggregate TCP approach in detail. Section 4 presents a detailed performance evaluation of MPAT using experiments in controlled settings as well as over the Internet, and compares MPAT with proposed techniques. Section 5 concludes the paper, detailing our key findings.

2 Existing Approaches

One class of approaches (e.g. pTCP [32]) that provide network QoS in a best-effort network, try to open extra simultaneous TCP connections, instead of a single connection, to give more bandwidth to an application. By contrast, MPAT does not create extra TCP connections. It tries to provide relative QoS among only the TCP flows that the applications normally create. The resulting behavior of schemes like pTCP is clearly not desirable, and in fact CM[2] proposes to explicitly forbid this kind of behavior by using one active AIMD loop (i.e. one congestion “probe”) for multiple flows between two given endpoints. Further, this approach does not scale to higher performance ratios between flows, since a large number of flows active at a bottleneck lead to significant unfairness in TCP [38]. Of course, the scalability of both MPAT and pTCP is limited by the delay-bandwidth product of the aggregated flows. But pTCP would reach this limit much before MPAT. As an example, we later show in Section 4.5 (using experiments on the real Internet) that MPAT can give 95 times more bandwidth to one application over another using only three TCP connections. On the other hand, schemes like pTCP would need to open 95 parallel TCP connections for this.

Key et. al. [5] propose to set different ‘flow control’ limits for various flows in order to provide different bandwidth to each flow. Please note that such schemes can only give lower bandwidth to a flow (as compared to its fair share), thus violating the *utilization* condition described in Section 1.

TCP Nice [26] is a way of providing only two-level QoS, while our scheme can provide multi-level QoS. TCP Nice can only give lower bandwidth to the background traffic. It cannot give higher bandwidth to the more critical traffic. We want to design a system that can give more bandwidth to one TCP flow at the expense of lower bandwidth to another flow, in such a way that the sum total of the bandwidth that the two flows get is same as the fair share of two TCP flows.

The idea of providing performance differentiation between

TCP flows sharing a path has been discussed in the context of aggregate congestion management approaches [10, 7] like TCP Session [1, 4], Congestion Manager (CM) [2, 13], TCP Trunking [21]. A TCP trunk is a TCP connection between two routers, which is shared by multiple TCP flows active between the two routers. The end-to-end flows terminate at the end points of the trunk, where data from all the flows is buffered. The bandwidth available on the trunk is the bandwidth available on the TCP trunk, which evolves using the standard TCP AIMD[30, 31] algorithm. The data from various TCP flows can then be sent over the trunk according to any chosen scheduling policy. CM[2] is a congestion management scheme that provides unified congestion management to TCP as well as non-TCP flows, decouples congestion control from reliability, and ensures that end-points of a connection cannot hog more than their fair share of network bandwidth by opening multiple connections between them. CM defines an aggregate (termed “macroflow”) as a set of TCP flows between a source and a destination host (termed “microflows”). Both CM and TCP Session keep one AIMD bandwidth estimation loop active per aggregate, and hence the bandwidth available to the aggregate is that entitled to one TCP flow. Microflows can share this bandwidth according to any chosen scheduling policy. Thus, in case of a logical overlay link, flows between the endpoints of the link could be aggregated into a CM macroflow or a TCP trunk, and performance differentiation could be provided between them.

However, by virtue of using one AIMD loop per trunk or per CM macroflow, both TCP Session and CM, in their current form, do not satisfy the *utilization* requirement. If N flows constitute an aggregate (trunk or macroflow), and the aggregate shares a bottleneck link with M background flows, then with equal sharing within the aggregate, the share of the bottleneck link bandwidth received by the aggregate’s flows is $1/(N(M + 1))$. If the N flows compete for bandwidth like standard TCP flows, each of the flows would be entitled to a share of $1/(M + N)$. This has also been mentioned in Chapter 7 of [1].

The above problem could be addressed by using a variant of the AIMD loop that acts like N TCP flows, as proposed in [1]. When a TCP flow has a bottleneck link where the loss probability p characterizes the congestion state at the link, TCP’s AIMD algorithm allocates the flow a bandwidth given by $B = K/(RTT * \text{sqrt}(p))$ [19], where K is proportional to $\sqrt{\frac{\alpha}{\beta} * (1 - \frac{\beta}{2})}$. To achieve performance differentiation among these flows, one possible approach is to play with the parameters α and β in the AIMD algorithm. This is the approach used by mulTCP [3]. The idea behind mulTCP is that one flow should be able to get N times more bandwidth than a standard TCP flow by choosing $\alpha = N/2$ and $\beta = 1/2N$.

Thus, the congestion window increases by N (as opposed to 1) when a congestion window worth of packets is successfully acknowledged. Losses leading to fast retransmit cause the window to be cut by $(1 - \beta)$, and losses leading to a timeout cut down the window to 1. Analytically, a mulTCP flow achieves a bandwidth N times that of a standard TCP flow experiencing the same loss rate. It is thus possible for CM or TCP trunking to use one mulTCP AIMD loop per aggregate to address the utilization problem. An equivalent TFRC [23] variant of mulTCP exists, where the throughput equation of mulTCP can be used in conjunction with a TFRC rate adjustment loop. A TFRC variant of mulTCP is used by OverQoS[9] to estimate the bandwidth entitled to N TCP flows on a logical overlay link.

Note however, that the loss process induced by a single mulTCP flow is quite different from that generated by N independent TCP flows. Hence, it is not clear if a mulTCP flow, especially for large N , would continue to behave like N TCP flows in the network. Also note that the ‘amplitude’ of mulTCP’s AIMD control loop increases with N , leading to an increasingly unstable controller as N grows. This is in contrast to N control loops of N independent TCP flows, where each has a small amplitude and thus tends to be more stable. This also has implications for a TFRC variant of mulTCP, since for large N , the analytically derived bandwidth equation of mulTCP may not represent the actual bandwidth consumed by a mulTCP flow.

Further, it has been noted in [3] that a mulTCP flow cannot act exactly like N independent TCP flows because timeout losses force the entire mulTCP window to be cut down to 1, whereas with N independent TCP flows, such a loss would cut down only one TCP connection’s window to 1. In [3], it is shown that this limits the value of N for which a mulTCP flow can achieve as much bandwidth as N independent TCP flows (the recommended value for N is 4 [8]).

The above discussion indicates that QoS approaches based on mulTCP may not meet the fairness and utilization requirements for large N , in turn violating the scalability requirement. In the remaining part of this section, we experimentally investigate this hypothesis.

We conducted experiments on the Emulab network testbed at University of Utah [11] over an emulated link of bandwidth 90Mbps and RTT 50msec. We ran a mulTCP flow along with 10 standard TCP flows through the same bottleneck link. To create a mulTCP connection equivalent to N TCP flows, we set the additive increase factor (α) and multiplicative decrease factor (β) such that $\frac{\alpha}{\beta} = N^2$ [37, 36, 32]. We performed all our experiments with two sets of parameters: $\alpha = N$ (i.e. $\beta = 1/N$), and $\alpha = \sqrt{N}$ (i.e. $\beta = N^{-1.5}$). We used a TCP-SACK based implementation of mulTCP.

Figure 1(a) shows how the achieved bandwidth ratio be-

tween multTCP and a standard background TCP flow varies with N . Figures 1(b) and 1(c) show the corresponding number of fast retransmits and timeouts seen on the multTCP flow. Note that for large N , due to multiple losses within the same window, fast retransmit does not trigger, causing a large number of timeout-induced window resets. The result is that multTCP's achieved bandwidth ratio flattens out for $N > 20$ for $\alpha = \sqrt{N}$, and actually falls for $\alpha = N$.


The 45-degree line in figure 1(a) represents the ideal achieved bandwidth ratio for a given choice of N . Note that for small values of N , multTCP's achieved bandwidth ratio lies above this graph, indicating that multTCP is more aggressive than N TCP flows to the background TCP flow.

Figure 2(a) shows how the mean and variance in the bandwidth achieved by multTCP varies with N . Note that the variance increases with N , till it eventually flattens out together with the bandwidth. This increased variance is a manifestation of the increasing instability of the modified AIMD controller. Note that even though multTCP's achieved bandwidth flattens out for large N , figure 2(b) shows that the background TCP flow is also not able to get the remaining bandwidth, as its throughput also flattens out. The result, as shown in figure 2(c), is a loss of utilization of the bottleneck link. The reason for this behavior is also an increased number of timeouts both for multTCP and for background TCP traffic, induced by multTCP's aggressive control loop for large N .

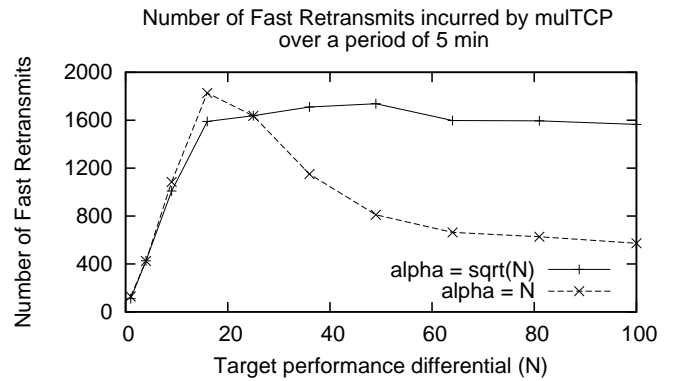
To analyse the behavior of multiple virtual overlay links using multTCP, we measured the bandwidth distribution achieved between two multTCP flows with $N = N_1$ and $N = N_2$. As shown in table 1, the achieved bandwidths are not in the ratio N_1/N_2 . This suggests that the interaction between two control loops, both with large amplitudes, does not lead to the ideal bandwidth distribution. The fairness of the bandwidth distribution tends to be much better with identical control loops, as used by standard TCP flows. Moreover, the utilization of bottleneck link (given by $B_1 + B_2$) keeps decreasing with larger values of N_1 and N_2 .

Finally, we study the behavior of a TFRC flow using multTCP, to verify if the slower responsiveness and averaging effects of TFRC mitigate some of multTCP's problems for large N . In this experiment, 5 multTCP TFRC flows run simultaneously with 5 background TCP flows. Each multTCP TFRC flow uses the multTCP bandwidth equation, and a measurement of the loss probability, to adjust the sending rate in response to loss events. Table 2 shows the average bandwidths achieved by a multTCP flow compared to a TCP flow and their ratio, for varying N . Note that the ratio grows much faster than N , which indicates that the multTCP TFRC flows aggressively take bandwidth away from background TCP flows as N increases.

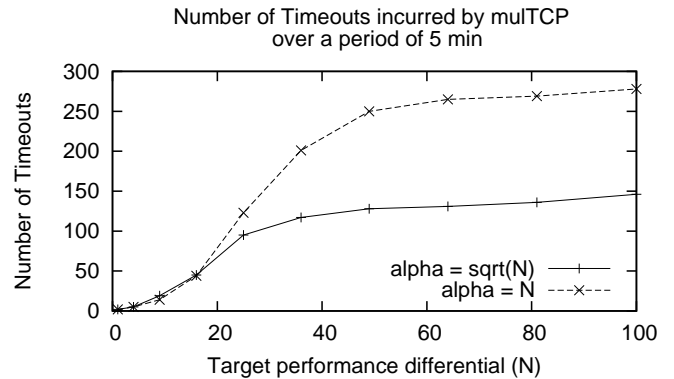
The above analysis shows that beyond small values of N ,



(a)

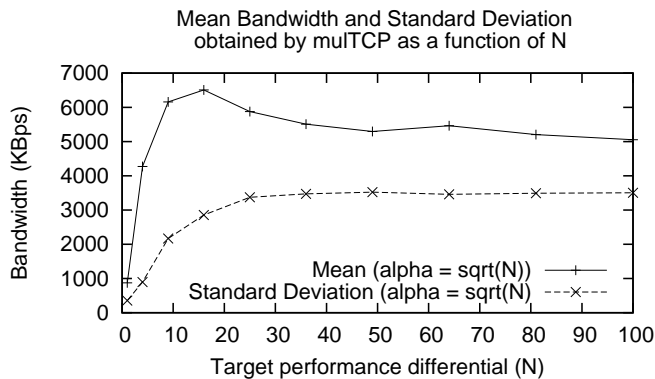


(b)




(c)


Figure 1: As N increases, due to multiple losses within the same window, fast retransmit does not trigger, causing a large number of timeouts, and hence limiting the scalability of the multTCP scheme.



(a)



(b)



(c)

Figure 2: As N increases, variance of a multTCP flow keeps increasing, and the total link utilization decreases.

N_1	N_2	B_1 (KBps)	B_2 (KBps)	B_1/B_2
4	1	3990	498	8.01
9	1	5519	392	14.07
16	1	6091	307	19.84
25	1	5907	280	21.09
9	4	4876	1795	2.72
16	4	5156	1320	3.91
16	9	4340	2833	1.53
25	4	5596	1171	4.78
25	9	4082	2618	1.56
25	16	2925	2264	1.29

Table 1: Table showing the bandwidth obtained by two multTCP flows running simultaneously through the same bottleneck link over Emulab network.

N	$TFRC$ (KBps)	TCP (KBps)	$TFRC/TCP$
5	1857	350	5.3
7	2078	166	12.5
8	2153	101	21.3
9	2196	63	34.8
10	2223	37	60
20	2230	22	101
30	2247	13	172
40	2254	7	322

Table 2: Table showing the bandwidth obtained by TFRC(+multTCP) connections along with standard TCP connection.

multTCP does not provide an adequate solution to the problem of creating an N -TCP abstraction. In particular, QoS schemes based on multTCP do not satisfy the scalability requirement. In the next section, we present an aggregate congestion management scheme that satisfies the fairness, utilization as well as scalability requirements mentioned in section 1.

3 QoS through aggregate congestion management

3.1 An illustration using two flows

Before we discuss how our MPAT scheme exactly works, let us start with a simple example. Consider the following scenario: we have two TCP flows running simultaneously between the same end-hosts, thus sharing the same bottleneck link and ex-

periencing similar delay. For ease of illustration, we refer to the first flow as the *red* flow and the second flow as the *blue* flow.¹

Our goal is to provide performance differentiation among these two flows under the following constraints: Firstly, we should not affect other background flows running in the network at the same time (in terms of bandwidth, delay, loss rate, etc). This is referred to as the *fairness* property in Section 1. Secondly, the sum total of bandwidth acquired by the red flow and the blue flow should be equal to the fair share of two standard TCP flows going through the same bottleneck link at the same time. This is referred to as the *utilization* property in Section 1. The reason this problem is hard is that the fair share of a TCP flow keeps changing dynamically with time, depending upon the number of background flows going through the bottleneck link, and the end-hosts do not have an explicit indication of the amount of such cross-traffic.

Suppose we want to apportion the available bandwidth² among these two flows in the ratio 4:1. Assume that the congestion windows for each of the two flows is 5. Thus, the network allows us to send 10 packets every Round Trip Time (RTT). In the case of standard TCP, we would have sent 5 red packets and 5 blue packets every RTT. But the network does not really care how many red and blue packets we actually send, as long as the total number of red and blue packets is 10. If we *aggregate* the congestion information of both these flows at the end-host (the sender side), we can transmit 8 red packets and 2 blue packets every RTT, thus giving four times more bandwidth to the red flow over the blue flow. This would also make sure that we do not hurt the background traffic at all.

The above step allows us to split the instantaneous bandwidth among the two flows in the ratio 4:1. But will we be able to achieve that over a long period of time? The answer is NO, and the reason is as follows. Since the probability of each packet getting dropped in the network is the same, the red flow would experience a higher loss rate than the blue flow. This would force the red flow to cut down its congestion window more often than the blue flow. In the long run, the total bandwidth acquired by the two flows would be much less than the fair share of two TCP flows, thus violating the *utilization* property.

To overcome this problem, the fundamental *invariant* that we try to maintain at all times is that *the loss rate experienced by each congestion window should be the same as in standard TCP*. Standard TCP sends data packets, and receives conges-

tion signals (either explicitly in terms of ECN, or implicitly in terms of duplicate acks, fast retransmissions, timeouts, etc) back from the receiver. If we had used standard TCP for each of the two flows, both the red window and the blue window would (on an average) experience equal number of congestion signals. In order to maintain this property, we first separate reliability from congestion control, as proposed in CM[2]. Next, we decouple the actual growth of congestion window from the identity of the flow whose packets advance the congestion window.

In the above example, when we get 8 red acks, we send three of these to the blue congestion window. In other words, we assign 5 acks to the red window, and 5 to the blue window (3 red + 2 blue). This ensures that each of the two congestion windows experiences similar loss rate, even though we can split the available bandwidth in the ratio 4:1. Please note that since we separate reliability from congestion control, it is the red flow that is responsible to maintain the reliability of each of the 8 red packets (in terms of buffering, retransmission, etc). When we get 8 red acks, we separate each of them into 'reliability ack' and 'congestion signal'. For the purpose of reliability, we send each of the 8 red acks to the red flow. But for the purpose of congestion control, we apply only 5 of these red acks to the red window, and apply rest 3 to the blue window. This ensures that the sum total of bandwidth acquired by both the flows is equal to the fair share of two TCP flows, thus satisfying the *utilization* property.

3.2 The general case

A source of N TCP flows sharing a bottleneck link is entitled to a total share of the link bandwidth given by the sum of the fair shares of each of the N TCP flows. The source should thus be able to arbitrate this total share among the N flows according to their performance requirements. Each TCP connection opened by an application has a corresponding control loop which uses the AIMD algorithm to adjust its congestion window in response to feedback from the network. The key idea behind our aggregate congestion management scheme is to *keep as many AIMD control loops active in the network as the number of TCP flows in the aggregate*, but to decouple application flows from their congestion windows. We call this scheme MPAT, for Multi-Probe Aggregate TCP, to emphasize the fact that we keep N AIMD control loops (i.e. "probes") active.

Thus, in an aggregate of N flows, the N application data streams are decoupled from the N congestion windows that are each evolving using the AIMD algorithm. The N AIMD control loops allow us to hold the fair share of N TCP flows, while an additional step of mapping application packets to congestion windows allows us to arbitrate the aggregate band-

¹Please note that in contrast to schemes like pTCP, we are not intentionally opening these flows in order to hog more network bandwidth. These are the flows that applications normally create.

²In this example, by 'available bandwidth', we refer to the fair share of two TCP flows.

width to provide the desired QoS. Since the identity of packets driving a given AIMD loop is irrelevant, this remapping is transparent to the network. Thus, the aggregate appears as N standard TCP flows to the network, providing appropriate fairness to any background TCP or TFRC traffic. Please note that this is different from all the existing schemes[2, 3, 21], which keep only one bandwidth estimation probe active in the network, and hence suffer from problems of scalability and fairness.

The following section describes MPAT in detail.

3.3 The MPAT Algorithm

Consider N TCP connections running simultaneously as part of an aggregate, sharing the same bottleneck link in the network. Let these flows be labeled as f_i , $1 \leq i \leq N$. Let C_i represent the congestion window of flow f_i . We introduce another variable A_i which denotes the MPAT window for flow f_i . Let C denote the aggregate congestion window, given by the sum of the congestion windows of all flows. Let x_i denote the fraction of the total bandwidth allocated to flow f_i , such that $\sum_i x_i = 1$. The shares x_i are derived from the performance requirement of the flows f_i , and could change dynamically with time.

Note that C represents the product of the bandwidth available to the aggregate and the round-trip delay (RTT) on the logical overlay link. While TCP would allocate C among the N flows roughly equally in every RTT, MPAT would allocate C in accordance with the performance requirements of the flows. In other words,

$$A_i = x_i * C \tag{1}$$

The actual number of packets that flow f_i is allowed to send every RTT is $\min(A_i, W_i)$, where W_i is the minimum of the sender and receiver window sizes for flow f_i . With standard TCP, flow f_i would be allowed to send $\min(C_i, W_i)$ packets every RTT.

C_i	Congestion Window for flow i
A_i	MPAT Window for flow i
C	$\sum_i C_i$
x_i	Bandwidth share of flow i
W_i	Minimum of sender and receiver window size for flow i

Table 3: Symbols and their meanings

Each connection i maintains a mapping of the sequence numbers of its packets to the identifier of the congestion window through which the packets were sent. We refer to this

mapping as $seqno2id_i()$. This mapping is maintained as a list ordered by increasing sequence number. For each congestion window i , we also maintain the inverse of this mapping, i.e. the list of packet sequence numbers sent through this window. We refer to this mapping as $id2seqno_i()$. This mapping is also maintained as a list ordered by increasing timestamps (i.e., the time at which the packet was sent). We use a variable per congestion window that keeps track of the number of outstanding packets that have been sent through that window.

3.3.1 Transmit Processing

Whenever connection f_i sends out a packet with sequence number s , it tries to find a connection (say j) with open congestion window. Congestion window j is said to be open if the number of outstanding packets sent through it is less than C_j . Note that the packets sent using congestion window j could belong to f_j , or to any other flow in the aggregate. In practice, j could be the same as i . Connection i then stores the mapping of seqno s to congestion window j in $seqno2id_i()$. We also store the inverse mapping in $id2seqno_j()$

3.3.2 Receive Processing

Our implementation is based on TCP-SACK. When we receive an acknowledgement for a packet with sequence number s , belonging to flow f_i , we use the mapping $seqno2id_i()$ to find the congestion window j through which the packet was sent. We then look at the inverse mapping $id2seqno_j()$ to find if the ack received is in sequence or not. We then apply the standard AIMD algorithm to congestion window j . Thus, if the ack is in sequence, we linearly increase C_j to $C_j + 1/C_j$. If it is a duplicate ack (in case of SACK, this would be the ack for a later packet in sequence), we enter the fast retransmit phase, halving C_j if we get three duplicate acks (in case of SACK, three out-of-sequence acks).

3.3.3 Mapping sequence numbers to congestion windows

The choice of the mapping of a sequence number to a congestion window is critical, since it affects the number of fast retransmit-induced window halvings that an MPAT aggregate receives over a given interval of time, and hence the total bandwidth gained by it. Recall that standard TCP-SACK considers multiple losses within a single congestion window as one loss event for fast retransmit, and hence cuts down its window by half only once. To understand how this affects MPAT, consider the following example: There are two TCP flows running simultaneously as part of an MPAT aggregate. Connection 1 has a window size of 20, and has packets 1 through 20 outstanding in the network. Suppose packets 8 and 9 get dropped. Since

both these packets belong to the same window, standard TCP-SACK would treat them as one signal for fast retransmit, and hence cut down its window only once.

MPAT could treat this situation as either one or two fast retransmit signals, depending upon the mapping $seqno2id_1()$. Consider the case when sequence numbers 8, 10-12 are mapped to C_1 , and sequence numbers 9, 13-15 are mapped to C_2 . The acks for packets 10-12 will be charged to C_1 , which will treat all three of them as duplicate acks for packet 8, and hence cut down its window by half. But the acks for packets 13-15 will be charged to C_2 , which will again treat them as duplicate acks for packet 9, and hence cut down its window by half. This example shows how two packet losses within the same window could lead to two fast retransmit signals. On the other hand, if both packets 8 and 9 were mapped to C_1 , this would lead to only one fast retransmit signal. In this case, flow 2 would treat acks for packets 13-15 as if they were received in sequence, and hence would *not* reduce its window.

We tried out three different algorithms to map the packet with sequence number s , belonging to flow f_i , to a congestion window :

- If sequence number $s - 1$ was sent through congestion window k , send the current packet s through window k if it is open. Else, randomly pick an open congestion window to send the packet. This algorithm turns out to be aggressive to the background flows since consecutive packet losses are charged to the same congestion window, and get absorbed into one loss event for fast retransmit, reducing the number of window reductions.
- If sequence number $s - 1$ was sent through congestion window k , randomly choose an open congestion window different from k . This scheme turns out to be more conservative than standard TCP, since it forcibly maps multiple losses on a flow to different congestion windows, causing each one of them to be cut down.
- Map s to an open congestion window picked uniformly at random. This option turns out to be a reasonable compromise between the two extremes. In fact, it turns out to be conservative when one of the flows has a very large share of the aggregate congestion window, since it still spreads out the losses of packets on that flow to multiple congestion windows. We use this option in our implementation.

Whenever any of the congestion windows in the aggregate changes, the aggregate congestion window also changes, and hence we must also update the MPAT windows A_i of all connections using equation 1. Note that this means that whenever the congestion window for one of the flows is cut down, the loss in available bandwidth gets distributed proportionally among the MPAT windows of all the flows. This has the effect

of reducing the inter-connection variance in throughput, as we shall see later in section 4.6.

3.3.4 TFRC variant

There exists a TFRC[23] variant of MPAT, in which instead of N AIMD control loops, we have N rate adjustment loops, each driven by the standard TCP throughput equation for one TCP flow. Note that loss probabilities will be measured by congestion window, i.e. over data sent using a given rate adjustment loop, irrespective of which application flow the packets came from. Loss events will be mapped to congestion windows in the same way as described above. The TFRC variant enjoys the same transparency properties as the AIMD version, in that each rate adjustment loop is oblivious to the identity of the packets driving it. As a result, the TFRC variant of MPAT appears as a set of N independent TFRC flows to the network. All of our experiments in this paper are with the AIMD version of MPAT.

3.4 Implementation

We have prototyped our proposed aggregate congestion management scheme in Daytona[28], a user-level TCP stack running on Linux. Daytona can be linked as a library into an application, and offers the same functions as that provided by the UNIX socket library. Daytona talks to the network using a raw IP socket, thus avoiding any processing in the operating system's TCP stack. All TCP processing is done at user-level, and pre-formed TCP packets are handed to the IP layer to be sent over a raw socket.

Our user-level implementation allowed us to run wide-area experiments on diverse locations in the Internet, as well as in controlled public network testbeds like Emulab [11], without requiring control over the operating system.

A key goal in MPAT is to ensure that we hold the fair share of N TCP flows in the process. Thus, each AIMD control loop should be kept running as long as there are packets available to drive it. Note that if a flow with a high share of the aggregate congestion window gets send or receive window limited, or has a period of inactivity, then it may leave some of its share of the aggregate congestion window unused. Left alone, the aggregate would lose this unused share to the background traffic. By using a work-conserving scheduler on the aggregate, however, we make sure that other flows in the aggregate take up such unused bandwidth, if they have data to send.

An implementation detail concerns the amount of memory allocated per socket corresponding to a connection. Due to an unequal apportionment of the total bandwidth among different connections, each connection needs a different amount of buffer memory at any given instant of time. Currently our

implementation statically partitions a large chunk of memory between connections in accordance with their performance requirements. In the future, our implementation would dynamically allocate socket memory, as proposed in [20].

4 Evaluation

In this section, we conduct extensive experimentation to evaluate the fairness, utilization and scalability properties of MPAT. We have conducted experiments both over the *real Internet* and in controlled settings under stable conditions, reverse traffic, and transient congestion. We built our own wide-area network testbed consisting of a network of nodes in diverse locations (including US, Europe and Asia). Our main goal in choosing these nodes is to test our system across wide-area links which we believe have losses. For this reason, we made sure that the whole path is not connected by Internet2 links (known to have very few losses). We changed the routing tables of our campus network to make sure that the path to our destination did not take any of the Internet2 links. We also did experiments on the Emulab network [11], which provides a more controlled setting for comparing our scheme with other approaches. All of our experiments used unconstrained send and receive windows on both ends of the TCP connections, so that flow control does not kick-in before congestion control. As background traffic we use long lived TCP connections and bursty web traffic. We had a Maximum Segment Size (MSS) of 1460 bytes in all our experiments.

4.1 Scalability

Using experiments on both the Emulab network and the real Internet, we found that the MPAT scheme can hold the fair share of 100 TCP flows running simultaneously as part of an aggregate, irrespective of how the aggregate bandwidth was apportioned among the aggregate's flows. The limits on scalability beyond 100 flows arise due to other factors like minimum amount of bandwidth needed per connection, bottleneck shifting from network to memory, etc.

As noted in section 2, mulTCP scales only up to 20-25 flows in the Emulab setting. In a real Internet experiment, we found mulTCP to scale upto 10 to 15 flows.

4.2 Performance differentiation

The MPAT scheme can be used to apportion the total available bandwidth among different flows in any desired ratio. Even though the total fair share allocated to the aggregate keeps changing with time, MPAT can maintain the target performance differential at all times. We could give a maximum of 95 times more bandwidth to one TCP flow over another.

The real limits to scalability for large performance differential arise from the fact that every connection needs to send a minimum of 1-2 packets every RTT [38]. In the future, we plan to change this to 1 packet every k RTTs, as proposed in [26].

As an example, Figure 3(a) shows the bandwidth allocated to five TCP flows running as part of an MPAT aggregate between Utah and our campus network on the east coast (RTT approximately 70 msec). There were five more long-running standard TCP flows in the background (in both directions). We tried to split bandwidth in the ratio 1:2:3:4:5. The graph is for a period of 5 min, with data samples taken every 500msec (approx 7 RTTs). The average amount of bandwidth that each of the five flows got was 135 KBps, 265 KBps, 395 KBps, 530 KBps and 640 KBps respectively. We can see that the MPAT scheme is very effective at maintaining the desired target differential among different flows at all times, irrespective of the total share of bandwidth that the aggregate is entitled to. The average bandwidth of a standard TCP flow running in background during this time period was 400 KBps. During the process of splitting bandwidth, the MPAT aggregate holds its total fair share of approximately 2000 KBps, as seen in Figure 3(b), thus satisfying the *utilization* criterion.

4.3 Fairness to background flows


As noted earlier, an MPAT aggregate is naturally fair to background TCP traffic since the remapping of application packets to congestion windows is transparent to the N AIMD control loops in the aggregate, which means that the network effectively see N standard TCP flows. Under the same experimental conditions as above, Figure 5(c) shows the ratio of total bandwidth occupied by the MPAT aggregate ($N = 16$) as compared to the total bandwidth of 10 standard TCP flows running in background.

4.4 Interaction between multiple MPAT flows

To analyse the behavior of multiple competing virtual overlay links using MPAT, we study the bandwidth achieved by multiple MPAT aggregates running simultaneously. Again, as noted earlier, the behavior of an MPAT aggregate with N flows is similar to N independent TCP flows running without any aggregation. This suggests that each MPAT aggregate should be able to get a share of the bottleneck bandwidth depending upon the number of flows in each aggregate, and the number of background TCP flows.


To verify this hypothesis, we ran 5 MPAT aggregates, each with a different value of N , together with 5 background TCP flows between Utah and our campus network on the east coast. The five MPAT aggregates use $N = 2, 4, 6, 8$ and 10. We

The MPAT scheme used to apportion total bandwidth in the ratio 1:2:3:4:5



(a) MPAT can apportion bandwidth among its flows, irrespective of the total bandwidth available.

Total bandwidth of the MPAT aggregate while splitting it in the ratio 1:2:3:4:5 among 5 flows



(b) The MPAT aggregate can hold its total fair share

Figure 3: MPAT can apportion bandwidth in the desired ratio, while holding its total fair share.

then replaced each aggregate with the same number of independent TCP flows. The experiment was repeated during different times of day. Figure 4 shows the bandwidth for each of the five MPAT aggregates, with data points sampled every 1sec. Tables 4(a) and 5(a) show the number of fast retransmits, (labelled # *Fast*), number of timeouts and bandwidth for each flow within the MPAT aggregate over a period of 390 seconds. Tables 4(b) and 5(b) show the corresponding data for the non-aggregated case. Note that in both cases, the achieved bandwidth and loss rates seen by the background TCP flows are similar. The bandwidth obtained by each aggregate is always proportional to the number of flows it had, and when the flows within an aggregate were run independently, the sum of bandwidths achieved by these flows was similar to the bandwidth achieved by the aggregate. This shows that MPAT is not a selfish scheme. Competing MPAT flows cooperate with each other, and with the background traffic. The experiment also illustrates that MPAT adequately satisfies the utilization property.

4.5 Adaptation to changing performance requirements

When the performance requirements of various flows within an aggregate change, our scheme simply needs to update the MPAT windows A_i for all flows using equation 1, and thus should be able to respond very quickly (typically 2-3 RTTs) to

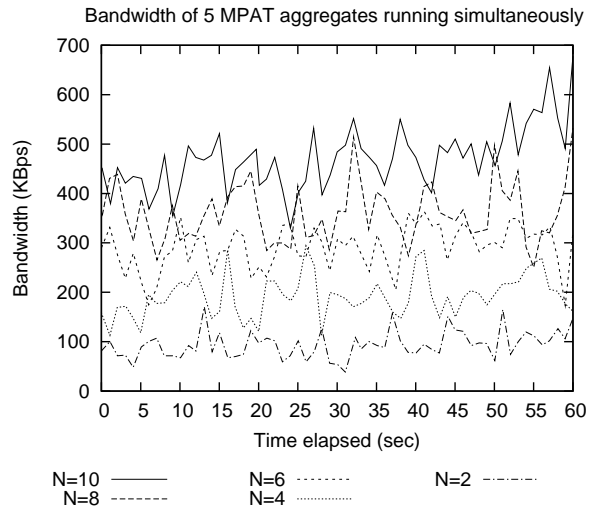


Figure 4: Competing MPAT aggregates cooperate with each other.

N	# Fast	# Timeouts	Bandwidth (KBps)
2 (MPAT)	511	250	51.4
4 (MPAT)	474	231	49.5
6 (MPAT)	499	235	50.3
8 (MPAT)	495	235	50.1
10 (MPAT)	502	236	49.8
5 (TCP)	401	258	48.6

(a) With Aggregation (day time)

N	# Fast	# Timeouts	Bandwidth (KBps)
2 (TCP)	498	218	56.1
4 (TCP)	479	234	58.5
6 (TCP)	481	215	58.9
8 (TCP)	481	216	57.1
10 (TCP)	490	218	57.2
5 (TCP)	497	213	56.3

(b) Without Aggregation (day time)

Table 4: Multiple MPAT aggregates running simultaneously cooperate with each other (day time).

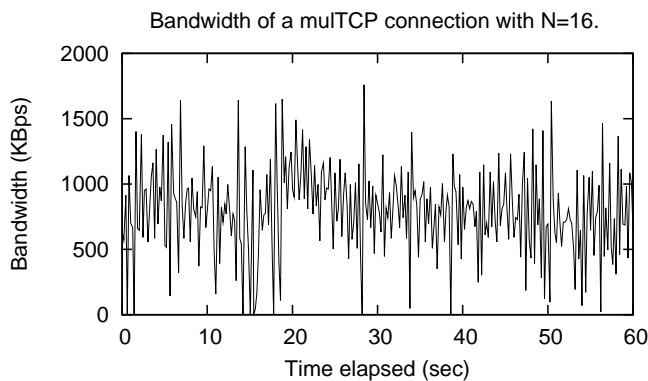
N	# Fast	# Timeouts	Bandwidth (KBps)
2 (MPAT)	565	44	97.6
4 (MPAT)	564	34	98.5
6 (MPAT)	555	39	98.1
8 (MPAT)	535	38	99.3
10 (MPAT)	537	41	97.9
5 (TCP)	577	41	93.6

(a) With Aggregation (night time)

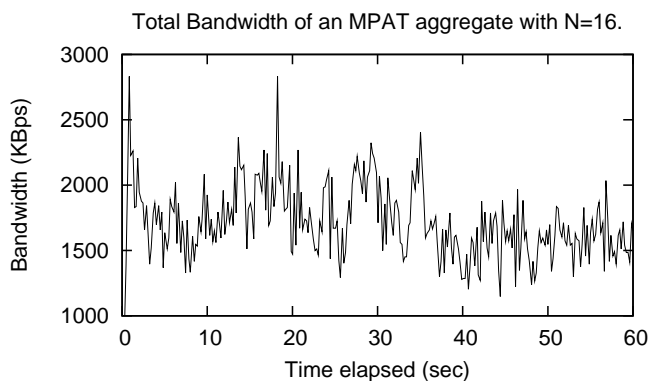
N	# Fast	# Timeouts	Bandwidth (KBps)
2 (TCP)	540	28	106.7
4 (TCP)	542	25	110.9
6 (TCP)	546	27	108.5
8 (TCP)	539	25	106.4
10 (TCP)	531	26	109.5
5 (TCP)	538	30	107.1

(b) Without Aggregation (night time)

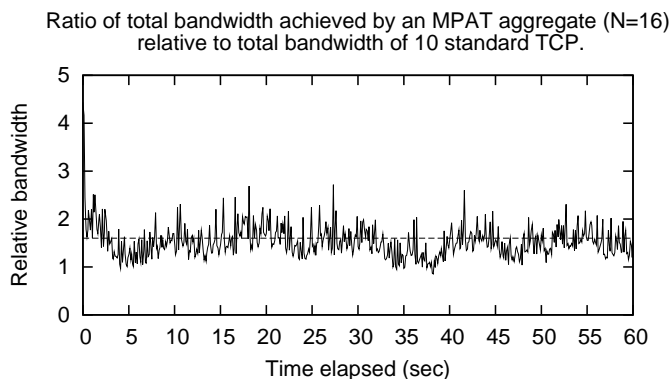
Table 5: Multiple MPAT aggregates running simultaneously cooperate with each other (night-time).



(a)



(b)



(c)

Figure 5: MPAT exhibits much lower variance due to sharing of losses, and is friendly to the background traffic.

dynamically changing performance requirements with time.

To test this claim, we ran an MPAT aggregate consisting of three flows between Utah and our campus network on the east coast, with 5 more long-running standard TCP flows in the background (in both directions). We changed the performance requirements every minute. Figure 6(a) shows how the absolute bandwidth for each of the three connections varies with time for a period of nine minutes. Data points have been sampled every 500 msec (approx 7 RTTs). Figure 6(b) shows relative bandwidth of the three flows at all times, along with the target performance differential. We can see that MPAT very quickly adapts to changing performance requirements. Note that during the time interval $t = 240$ to $t = 300$ sec, we were able to split bandwidth in the ratio 95:1.

4.6 Reduced variance

The total bandwidth of an MPAT aggregate is equivalent to the sum of N independent standard TCP flows, each of which uses an AIMD control loop. Multiplexing of these control loops smooths out the bandwidth variations within each AIMD loop, thus reducing the variance in throughput of an MPAT aggregate drastically. This is in contrast to multTCP that exhibits increasing variance with N .

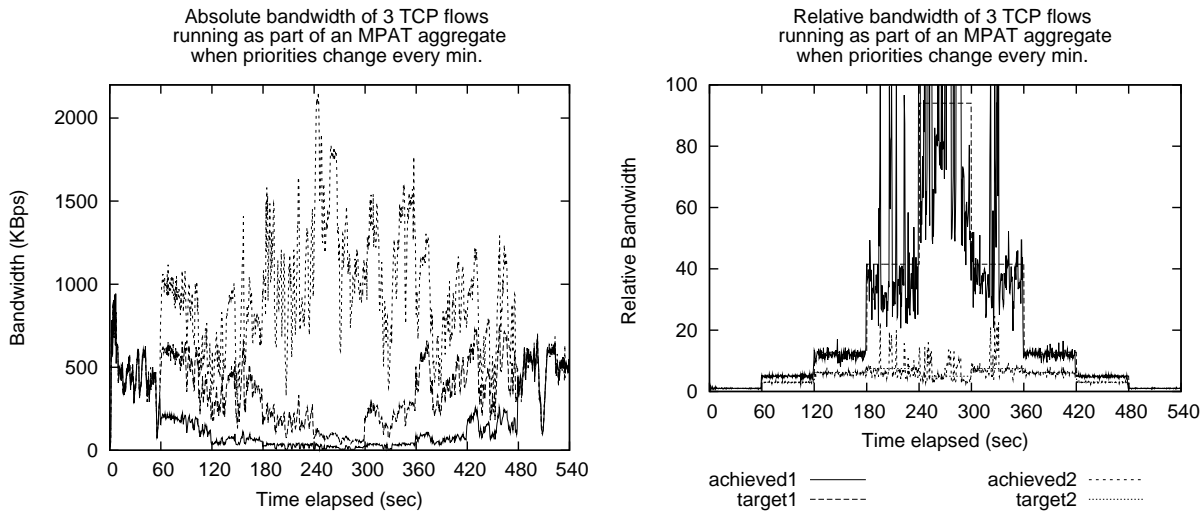
To demonstrate this, we ran an MPAT aggregate consisting of 16 flows between Utah and our campus network on the east coast. We had 10 long-running standard TCP flows running in the background. Figure 5(b) shows how the total bandwidth occupied by the MPAT aggregate varies with time. Figure 5(a) shows the bandwidth of a multTCP flow (with $N = 16$) under similar conditions. Data points have been sampled every 200msec (3 RTTs).

Comparing Figure 5(a) and Figure 5(b), we see that multTCP has a very high variance as compared to that of MPAT. Please note that the y-axis in Figure 5(b) is shown for the range 1000-3000 KBps, while that in Figure 5(a) is for the range 0-2000 KBps. This is because a multTCP connection with $N = 16$ gets much less throughput than the fair share of 16 TCP flows.

4.7 Adaptation to Transient Congestion

The MPAT scheme must be robust to transient congestion in the network, and more generally, to variations in the fair share of bandwidth available to the aggregate. With transient congestion, an MPAT aggregate must learn about a change in its total fair share using the AIMD algorithm. The change is apportioned among the aggregate's flows quickly through an adjustment of the MPAT windows A_i of all flows, ensuring that the desired bandwidth ratio is maintained at all times.

To study the responsiveness of MPAT to transient congestion, we conducted experiments on the Emulab network over



(a) This figure shows how the absolute bandwidth achieved by flows that are part of an MPAT aggregate ($N=3$) adapts itself when priorities keep changing dynamically with time.

(b) This figure shows how the relative bandwidth achieved by flows that are part of an MPAT aggregate ($N=3$) adapts itself when priorities keep changing dynamically with time.

Figure 6: MPAT Scheme can adapt itself very quickly to dynamically changing performance requirements.

an emulated link of bandwidth 90 Mbps and RTT 50msec. We ran an MPAT aggregate consisting of 10 flows with different performance requirements. Transient congestion was created by introducing constant bit rate (CBR) as well as TCP traffic in the background. We also used 10 long-running standard TCP flows in both directions. At $t=120$ sec, we introduce a CBR flow in the background with rate 30Mbps. As shown in Figure 7(b), MPAT cuts down its total bandwidth in about 4-6 RTTs (200-300 msec), while still apportioning bandwidth within the aggregate in the desired ratio at all times. At $t=240$ sec, we removed the background CBR traffic. As seen in figure 7(c), MPAT reclaims its fair share of the remaining bandwidth.

When background TCP traffic was introduced to create transient congestion, it takes about 15-20 (1 sec) RTTs for both MPAT and the background TCP flows to settle into their respective fair shares, as seen in figures 8(b) and 8(c). This is because the new TCP flows begin in slow-start mode, and are also adapting to their fair share together with MPAT. As earlier, MPAT always apportioned bandwidth within the aggregate in the desired proportion.


4.8 Bursty background traffic

Most of the results we described above are for long-running TCP flows in the background. We also tested our scheme with bursty web traffic running in background. We did this using


a *wget* loop downloading web pages in the background, together with MPAT. While bursty traffic increased the absolute number of retransmits and timeouts *proportionally* for all aggregates and the background traffic, we did not see any qualitative change in the results. In other words, MPAT exhibited the same scaling, fairness and utilization properties with bursty traffic.

5 Conclusions


The paper demonstrates for the first time the viability of providing differential services, at large scale, among a group of TCP flows that share the same bottleneck link. We demonstrate this through a range of experiments on the real Internet which show that an MPAT aggregate can hold its fair share of the bottleneck bandwidth while treating other flows fairly (either individual TCP flows or other MPAT aggregates). We also demonstrate that within an aggregate, MPAT allows substantial differentiation between flows (up to 95:1). This suggests that MPAT is therefore an appropriate candidate technology for broader overlay QoS services, both because MPAT provides the requisite differentiation and scalability, and because MPAT can co-exist with other TCP flows and with other MPAT aggregates. This result opens the door to experimentation with more easily deployable network QoS schemes such




(a) This figure shows how the total bandwidth achieved by an MPAT aggregate (N=16) adapts itself with additional UDP traffic c introduced at t=120 and then removed at t=240.



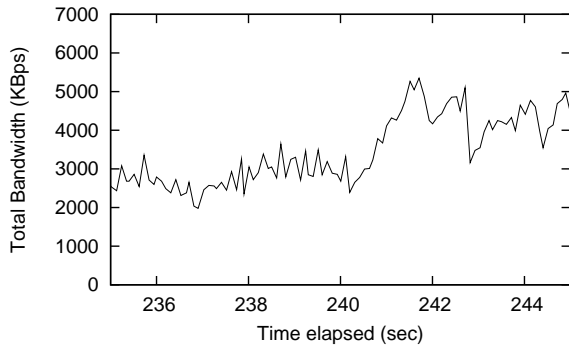
(a) This figure shows how the total bandwidth achieved by an MPAT aggregate (N=16) adapts itself with additional TCP traffic c introduced at t=120 and then removed at t=240.




(b) This figure shows how the total bandwidth achieved by an MPAT aggregate (N=16) decreases with additional background UDP traffic c introduced at t=120.



(b) This figure shows how the total bandwidth achieved by an MPAT aggregate (N=16) adapts itself with additional background TCP traffic c introduced at t=120.



(c) This figure shows how the total bandwidth achieved by an MPAT aggregate (N=16) increases with additional background UDP traffic c removed at t=240.



(c) This figure shows how the total bandwidth achieved by an MPAT aggregate (N=16) adapts itself with additional TCP background traffic c removed at t=240.

Figure 7: Behavior of MPAT Scheme with transient congestion induced by long-running UDP flows.

Figure 8: Behavior of MPAT Scheme with transient congestion induced by long-running TCP flows.

as the overlay.

Having said that, this paper does not make the broader conclusion that overlay QoS works. In order to do that, we must understand the end-to-end behavior of flows (TCP and otherwise) over a multihop overlay. This means among other things that we must understand the interactions between the individual overlay hops (be they TCP or TFRC) of the same overlay, between aggregates on different overlays, and between all of these and the end-to-end flow control of the TCP connections running over the overlay. It is also important to understand how an overlay QoS could be combined with RON-type overlay functionality used to actually enhance performance (not just provide differentiation).

There may also be other uses for MPAT aggregation. For instance, in the back end systems of a web service data center, MPAT could be used to provide differential service among different kinds of customers (gold, silver, bronze). This is possible in part because MPAT only needs to be deployed at the sending end of a TCP connection. For this to work, however, the TCP flows must share a bottleneck. This in turns requires that the server implementing MPAT can determine which flows, if any, share the bottleneck. These possible uses of MPAT provide exciting avenues for future research.

References

- [1] V. Padmanabhan, *Addressing the Challenges of Web Data Transport*, Ph.D. Thesis, University of California at Berkeley, Sept 1998. <http://www.cs.berkeley.edu/padmanab/phd-thesis.html>
- [2] H. Balakrishnan and H. S. Rahul and S. Seshan, *An Integrated Congestion Management Architecture for Internet Hosts*, in Proceedings of ACM SIGCOMM, pages 175-187, Sept 1999.
- [3] J. Crowcroft and P. Oechslin, *Differentiated End-to-End Internet Services using a Weighted Proportional Fair Sharing TCP*, ACM Computer Communication Review, 28(3):53-69, July 1998.
- [4] V. Padmanabhan and R. Katz, *Addressing the Challenges of Web Data Transport*, Unpublished, Jan 1998. <http://www.research.microsoft.com/padmanab/papers/thesis-paper.pdf>
- [5] P. Key, L. Massoulie and B. Wang, *Emulating Low-priority Transport at the Application Layer: A Background Transfer Service*, in Proceedings of Sigmetrics, New York City, June, 2004
- [6] D. Andersen and H. Balakrishnan and M. Kaashoek, and R. Morris, *Resilient Overlay Networks*, in Proceedings of 18th ACM SOSP, Oct 2001.
- [7] V. Padmanabhan, *Coordinating Congestion Management and Bandwidth Sharing for Heterogeneous Data Streams*, in Proceedings of NOSSDAV'99, Basking Ridge, NJ.
- [8] P. Gevros and F. Risso and P. Kirstein, *Analysis of a method for differential TCP service*, in Proceedings of IEEE GLOBECOM, pages 1699-1708, Dec 1999.
- [9] L. Subramanian and I. Stoica and H. Balakrishnan and R. Katz, *OverQoS: An Overlay Based Architecture for Enhancing Internet QoS*, Accepted for publication in First Symposium on Networked Systems Design and Implementation (NSDI), March 2004.
- [10] J. Touch, *TCP Control Block Interdependence*, RFC-2140, April 1997.
- [11] Emulab test-bed. <http://www.emulab.net>
- [12] Planetlab test-bed. <http://www.planet-lab.org>
- [13] D. Andersen and D. Bansal and D. Curtis and S. Seshan and H. Balakrishnan, *System Support for Bandwidth Management and Content Adaptation in Internet Applications*, in Proceedings of the 4th Symposium on Operating Systems Design and Implementation, pages 213-226, Oct 2000.
- [14] 3rd Generation Partnership Project 2 (3GPP2), *Quality of Service: Stage 1 Requirements*, 3GPP2 S.R0035, www.3gpp2.org
- [15] A. Bestavros and O. Hartmann, *Aggregating Congestion Information Over Sequences of TCP Connections*, In BUCS-TR-1998-001, Boston University, Dec 1998.
- [16] R. Braden and D. Clark and S. Shenker, *Integrated services in the Internet architecture: an overview*, RFC 1633, June 1994.
- [17] S. Blake and D. Black and M. Carlson and E. Davies and Z. Wang and W. Weiss, *An architecture for differentiated services*, RFC 2475, Oct 1998.
- [18] A. Akella and S. Seshan and H. Balakrishnan, *The Impact of False Sharing on Shared Congestion Management*, in Proceedings of ACM SIGCOMM Computer Communication Review, Jan 2002.
- [19] J. Padhye and J. Firoiu and J. Kurose, *Modelling TCP Throughput: A Simple Model and its Empirical Validation*, in Proceedings of ACM SIGCOMM, pages 303-314, Aug 1998.
- [20] J. Semke and J. Mahadavi and M. Mathis, *Automatic TCP Buffer Tuning*, in Proceedings of ACM SIGCOMM, pages 315-323, Aug 1998.
- [21] H.T. Kung and S.Y. Wang, *TCP Trunking: Design, Implementation and Performance*, in Proceedings of 7th International Conference on Network Protocols (ICNP'99), pages 222-231, 1999.

- [22] S. Savage, *Sting: a TCP-based network measurement tool*, in Proceedings of USENIX Symposium on Internet Technologies and Systems (USITS), Oct 1999.
- [23] S. Floyd and M. Handley and J. Padhye and J. Widmer, *Equation-based Congestion Control for Unicast Applications*, in Proceedings of ACM SIGCOMM, pages 43-56, Aug 2000.
- [24] P. Pradhan and T. Chiueh and A. Neogi, *Aggregate TCP congestion control using multiple network probing*, in Proceedings of ICDCS 2000, pages 30-37, April 2000.
- [25] H. Balakrishnan and V. Padmanabhan and S. Seshan and M. Stemm and R. Katz, *TCP behavior of a busy web server*, in Proceedings of IEEE INFOCOM, pages 252-262, March 1998.
- [26] A. Venkataramani and R. Kokku and Mike Dahlin, *TCP NICE: A mechanism for background transfers*, in Proceedings of the 5th Symposium on Operating Systems Design and Implementation, pages 329-344, Dec 2002.
- [27] C. Jin, D.X. Wei and S.H. Low, *FAST TCP: Motivation, Architecture, Algorithms, Performance*, to be published in Proceedings of IEEE Infocom, March 2004.
- [28] *The Daytona User-Level TCP Stack*, Reference hidden for anonymity.
- [29] L. Zhang and S. Deering and D. Estrin and S. Shenker and D. Zappala, *RSVP: A new resource ReSerVation Protocol*, IEEE Network, 7:8-18, Sept 1993.
- [30] D. Chiu and R. Jain, *Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks*, Journal of Computer Networks and ISDN Systems, 17(1):1-14, June 1989.
- [31] V. Jacobson, *Congestion Avoidance and Control*, in Proceedings of ACM SIGCOMM, pages 314-329, Aug 1988.
- [32] H. Y. Hsieh and K.H. Kim and R. Sivakumar, *On Achieving Weighted Service Differentiation: an End-to-end Perspective*, in Proceedings of IEEE IWQoS, June 2003.
- [33] P. Gevros, *Internet Service Differentiation using Transport Options: the case for policy-aware congestion control*, in ACM Workshop on Revisiting IP QoS (RIPQoS), Aug 2003.
- [34] B. Davie, *Deployment experience with Differentiated Services*, in ACM Workshop on Revisiting IP QoS (RIPQoS), Aug 2003.
- [35] P. Hurley and J.-Y. L. Baudec and P. Thiran and M. Kara, *ABE: Providing a Low-Delay Service with Best Effort*, IEEE Network Magazine, Special Issue on Control of Best-Effort Traffic, 15(3):60-69, May 2001.
- [36] C. Dovrolis and P. Ramanathan, *A case for Relative Differentiated Services and the Proportional Differentiation Model*, IEEE Network, 13(5):26-34, Sept 1999.
- [37] T. Nandagopal and K.W. Lee and J. R. Li and V. Bharghavan, *Scalable Service Differentiation Using Purely End-to-End Mechanisms: Features and Limitations*, in Proceedings of IEEE IWQoS, June 2000.
- [38] D. Lin and R. Morris, *Dynamics of Random Early Detection*, in Proceedings of ACM SIGCOMM, pages 127-137, Sept 1997.