

# Performance Optimization of TCP/IP over Asymmetric Wired and Wireless Links

Dibyendu Shekhar, Hua Qin, Shivkumar Kalyanaraman, Kalyan Kidambi

Department of Electrical, Computer and Systems Engineering  
Rensselaer Polytechnic Institute, 110, 8<sup>th</sup> Street, Troy, NY 12180-3590, USA

Ph: +1 518 276 8979, Email: [shekhd@rpi.edu](mailto:shekhd@rpi.edu), [qinh@rpi.edu](mailto:qinh@rpi.edu), [shivkuma@ecse.rpi.edu](mailto:shivkuma@ecse.rpi.edu)

## ABSTRACT

In this paper, we analyze TCP performance in asymmetric networks, where the throughput significantly depends on the reverse direction and packet loss. We introduce a new operational model called the “AMP model” which, with an understanding of TCP dynamics and buffer management techniques, explains the performance effects observed. We apply this model to guide design improvements in buffer management (ack-regulation) and scheduling schemes to achieve performance improvements of an order of magnitude or more. The schemes have been tested through simulation and implementation, and the module driven improvements observed are better than those proposed in earlier literature.

## 1 INTRODUCTION

Transport Control Protocol (TCP) is widely used in Internet communications. However, the TCP/IP performance over wired and wireless networks is significantly diminished due to the effects of asymmetry. Figure 1 shows a network topology with wired and wireless asymmetric channels: the forward channel from broadband networks or wireless networks ranges in speeds is much higher than the reverse channel does. This bandwidth asymmetry has a profound impact on TCP performance because the TCP throughput is regulated (and limited) by the flow of acknowledgements (called “ack clock”) [1]. The more asymmetric the bandwidth, the more severe impact will have on TCP performance.

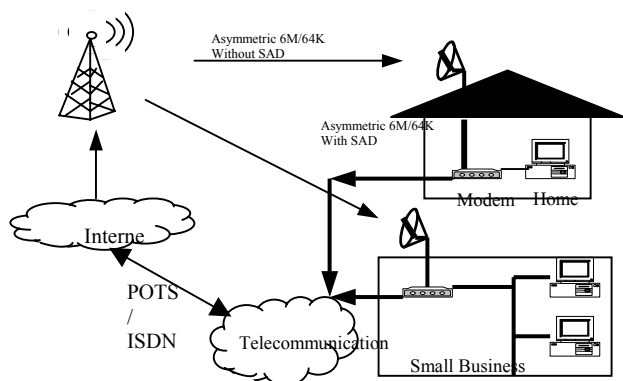


Figure 1: Network Topology with Asymmetric Links

The impact of asymmetry on TCP/IP performance has been characterized by Lakshman [2] using an index called *normalized asymmetry ratio* ( $k$ ), which is the

ratio of raw bandwidths on both directions to the ratio of packet sizes in both directions. When  $k$  is greater than one, the TCP throughput on the forward channel is restricted to a maximum of (forward channel bandwidth) /  $k$ . Other factors like bi-directional traffic (e.g.: downloading a web page while sending an email) or protocol overhead (e.g.: PPPoE, PPTP, L2TP, ATM etc) will increase  $k$  and further aggravate the asymmetry problem.

Balakrishnan [3], Keshav [4] and the PILC group [5] independently show that performance can be substantially increased by making two key changes: simply suppressing acknowledgements on the reverse channel (*ack-filtering*), and regenerating them after the reverse link has been traversed (*ack-reconstruction*). Balakrishnan also apply the same techniques to address types of asymmetry other than bandwidth asymmetry including asymmetry in delay, loss rates etc. We classify these techniques, as “*ack-regulation techniques*” since they are very different from mechanisms like Random Early Drop (RED) because the latter are designed to drop packets, not acks, and are used to signal congestion to TCP, not alleviate asymmetry problems. Our work focuses on generalizing these studies into a model called the *AMP model*, and making specific design improvements (Smart Ack Dropper: SAD) for performance optimization based upon this model.

## 2 THE AMP MODEL

We extend the notion of normalized bandwidth ratio ( $k$ ) [2] to form a simple model called the AMP model. The model is described as follows:

- Assume that, over an observation period, the fraction of the reverse link allocated to ack traffic is saturated at  $A$  acks/second, and
- Each ack generates  $M$  packets on the average ( $M =$  multiplicative factor), due to the effects of TCP dynamics and ack-regulation schemes working together, and
- The average size of packets is  $P$  bits/packet.

The forward link capacity allocated to packets corresponding to these acks is  $F$ bits/s. And the forward link throughput over the observation period is limited to  $\text{Min}(F, A * M * P)$  bits/second.

Observe that Lakshman’s normalized bandwidth ratio “ $k$ ” characterizes an absolute upper bound on the

forward throughput based upon link bandwidths and packet vs. ack sizes. The AMP expression, on the other hand, is an *operational bound* achievable with the asymmetric channel augmented with ack-regulation schemes. In particular, if the ack-regulation components can achieve a maximum average multiplicative factor of  $M$ , then the AMP expression is tighter bound on the achievable forward throughput. Moreover, it also accounts for scheduling allocations to acks and corresponding packets and vice versa (in the definition of  $F$  and  $A$ ), and can hence be useful in understanding effects of bi-directional traffic.

For example, if the reverse link speed is 64kbps, the forward link speed is 8Mbps, packet size is 1000 bytes, ack size = 40 bytes, TCP is in its slow start phase (generating two packets per ack) and no ack-regulation schemes are used. Then ack rate ( $A$ ) would now be 200 acks/second, the multiplicative factor  $M=2$ , and  $P = 8000$  bits, and the maximum throughput limit on the forward link =  $\text{Min}(8 \text{ Mbps}, 3.2 \text{ Mbps}) = 3.2 \text{ Mbps}$ . Note that *the number of TCP flows sharing the asymmetric link is immaterial in this model* (it is captured in  $M$ ). If TCP sources in particularly are all in congestion avoidance (and not in slow start), then  $M$  is closer to 1, and limit is even lower (about 1.6 Mbps)!

The AMP model is also useful in analyzing the bi-directional traffic case, especially when link-sharing schemes like CBQ [6] are deployed. Specifically, assume that in each direction, there are two classes (queues) served by CBQ: one for packets and one for acks. Further assume that packets on the forward link get a fraction  $f$  ( $1 > f > 0$ ) of the forward link capacity ( $C_f$ ), and on the reverse link get a fraction  $g$  ( $1 > g > 0$ ) of the reverse link capacity ( $C_r$ ). Acks, therefore get fractions  $1-f$  and  $1-g$  of capacity forward and reverse links, respectively. Assume that  $P_{\text{ack}}$  is the size of acks in bits/ack and that the reverse channel (both packets and acks) is saturated at their respective scheduling shares. The bounds on the maximum rates of packets and acks in both directions are given in (Table 1):

	Max Packet Rate (pkts/s)	Max Ack Rate (acks/s)
Reverse Channel	$G * C_r / P$ (saturated)	$(1-g) * C_r / \text{Pack}$ (saturated)
Forward Channel	$\text{Min}[(1-g) C_r / P, f * C_f / P]$	$\text{Min}[f * C_f / P_{\text{ack}}, g * C_r / P]$

Table 1: AMP model in bi-directional traffic

Observe that  $f$  (the link fraction allocated to packets on the forward link) should be chosen to be large, and can be as large as 99% for pure TCP/IP bi-directional traffic. The choice of  $g$  (the link fraction allocated to packets on the reverse link) is a tradeoff between performance and policy considerations. We will use this model to design ack regulation and scheduling policies.

### 3 THE SMART-ACK DROPPING (SAD) AND ACK REGENERATION POLICY

The Smart-ack dropping technique (SAD) is a simple extension of concepts developed by Balakrishnan [3]

and Srinivasan [4], i.e. to suppress as many acks in the reverse channel as possible because acks are cumulative. Balakrishnan's "ack-filtering" technique involves checking the entire queue upon the arrival of a new ack to remove earlier acks for the same connection. The goal is partially to free some space in the queue for other data packets and acks, and partially to compress the ack information. Srinivasan [4] independently proposes an "ack-collapsing" technique where all acks are queued, but at the transmission opportunity, the latest ack is sent and all others are dropped.

SAD is close to Srinivasan's "ack-collapsing" in concept, but uses minimal per-flow state to avoid enqueueing acks, which are going to be dropped anyway. This way, we require only a buffer of size  $N$  (acks) where  $N$  is the number of active flows (assuming separate packet buffers). In particular, SAD works as follows (Fig. 2):

We assume a FIFO queue for acks (and optionally packets). The per-flow information (stored as a hash table) includes a single bit, which indicates if an ack of that flow exists in the queue, and the latest ack number seen from that flow. When an ack arrives and the per-flow bit is zero, then we enqueue the ack, set the per-flow bit and copy the ack sequence number into the table. If the per-flow bit is already set upon arrival of a non-duplicate ack, then we update the per-flow ack number information to this value, and drop the ack. When an ack is dequeued for transmission, the latest value of ack number from the table is copied to the header (header checksum adjusted), and the bit in the table is cleared. Ack processing is  $O(1)$ , since we search the hash table and not the FIFO queue. This scheme can be extended to account for duplicate acks and SACKs [7].

A detailed analysis of the scheme applying the AMP model is also described in [7]. Specifically, if TCP is in congestion avoidance, the throughput is almost constant at  $W/\text{RTT}$ . The impact of SAD is to use the reverse channel capacity of  $A$  acks/s to support a maximum forward rate of  $\text{Min}[M * A, W/\text{RTT}]$  packets/s. When TCP is in slow start, each ack reaching the TCP source results in  $(M + 1)$  segments being sent where  $m$  is the number of acks suppressed by SAD. In fact, if the TCP source remains in slow start phase, the AMP model predicts that SAD can compensate for arbitrary degrees of asymmetry (i.e. any finite  $k$ ). However, since most TCPs reach congestion avoidance within a few RTTs, the multiplicative factor saturates at the value of (near-constant)  $M$  seen during the congestion avoidance.

SAD also reduces the probability of negative interactions with RTT estimation (which occurs in Balakrishnan et al's scheme) because acks are sent out in a timely manner to avoid spurious timeouts, and queueing delay is bounded by the time between enqueue and dequeue of any particular ack.

The suppression of acks leads to burstiness in the forward direction. This burstiness can be alleviated by using an *acknowledgement regeneration/reconstruction*

(AR) technique at the end of the reverse link as suggested by Keshav [4] and Balakrishnan [3]. AR is a method where the acks dropped by the SAD are regenerated. The ack coming to Ack-Regenerator carries the required information like the number of acks dropped and connection identifier encoded in a special field within the ack. To achieve this association, a per-connection table is maintained at the regenerator that maintains the information about the latest acked sequence number belonging to each connection. And a *regeneration factor* is introduced which denotes the number of acks generated in response to one ack.




Figure 2: Illustration of SmartAckDropper(SAD)

#### 4 PERFORMANCE EVALUATION

In this section, we use this actual model to design our experimentation architecture on "ns" simulation and VxWorks implementation. Finally, we present the results of simulation and implementation.

Figure 1 shows a general network topology, which illustrates bandwidth asymmetry. In our experiment, more specially, we study downstream channel ranges in speed from 6Mbps-8Mbps, and upstream channel ranges in speed from 64Kbps-768Kbps.

##### 4.1 Network Simulation

We use network simulation ns (version 2) as our experiment tool. We first present packet lossless simulation both in unidirectional and bi-directional traffic. Then, we present packet loss and delay simulation in unidirectional traffic.

###### 4.1.1 Lossless simulations

- Lossless Uni-directional TCP Flow

Table 1 shows the throughputs obtained for two configurations — TCP Reno and Reno with SAD but without any ack regulation schemes in uni-directional single flow — with different types of return channels. The simulation yielded the following results, which can be used as a baseline to evaluate the effectiveness of SAD. As shown, the asymmetry problem grows increasingly worse as the upstream channel becomes more constrained to a worst case of 19% utilization. However, these problems can be overcome almost entirely by using SAD. For instance, in the worst-case scenario (8 Mbps/64 kbps), throughput rose four-fold to 6.2 Mbps.

Additional throughput can be gained using AR techniques (Table 3). A regeneration factor (R) of R generates R acks for every ack suppressed by SAD. With a regeneration factor of 4, a maximum throughput

of 7.4 Mbps can be achieved even in the worst-case scenario.

Downstream/Upstream link (Mbps/Kbps)	TCP Reno without SAD	TCP Reno with SAD	K
	Throughput (Mbps)/Link Utilization	Throughput (Mbps)/Link Utilization	
8/64	1.52 / 19%	6.20 / 77.5%	5 <sup>1</sup>
8/128	3.01 / 37%	6.79 / 84.88%	2.5
8/256	5.98 / 74.75%	7.28 / 91%	1.25
8/640	7.46 / 93.25%	7.45 / 93.13%	0.5
8/784	7.46 / 93.25%	7.46 / 93.25%	0.41

Table 2: Uni-directional Single Flow Simulation without and with SmartAckDropper

Downstream/Upstream link (Mbps/Kbps)	Throughput (Mbps)	Regeneration factor
8/64	7.51	4
8/64	7.40	3
8/64	6.98	2
8/64	6.65	1

Table 3: SmartAckDropper and Ack Regenerator

Although this performance is the result of a trivial uni-directional flow, in reality one will find multiple flows, both uni-directional and bi-directional. We focus on the latter here for brevity. Please refer our detailed technical report for a full suite of simulation [7].

- Lossless Bi-directional TCP Flow

The problems associated with bi-directional transfers are unique and require a slightly different approach. Here, link utilization on *both* channels is key performance indicator. Balakrishnan [3] proposed that an acks-first scheduling policy be used. However, we find that if acks are given priority, the rate of packets on the reverse (low-speed) channel drops to unacceptably low values. Yet if packets are given priority, the effects of asymmetry lead to low packet throughput on the forward (high-speed) channel and beyond a point, this cannot be compensated by ack-regulation schemes (since the behavior is also dependent on TCP window dynamics).

Table 4 shows performance of a single FTP flow in each direction (baseline), with FIFO queuing at the upstream link. Here, the presence of bi-directional traffic reduces downstream throughput (and increases asymmetry) by almost a factor of 10. Downstream throughput worst case is 18.8 kbps, compared to 1.5 Mbps in the uni-directional model (Table 2). In the case of 384kbps upstream, downstream throughput is reduced to 1.57% available capacity. This drop can be attributed almost entirely to increased asymmetry and bidirectional issues, as there were no lost packets observed.

<sup>1</sup> (8 Mbps/64 kbps)/(1000 bytes/40 bytes) = 5

The cleanest way to improve performance in such cases is to set up one queue for acks and one for packets in both directions, serviced with a link-sharing scheduler (CBQ: Class-Based Queuing) rather than a priority scheduler. On the ack queue on the reverse direction, both SAD and AR are applied. CBQ assigns fractions, F and G respectively, to the packet queues on the reverse and forward links respectively. Based on analysis in section 2, we set F large — as large as 99.7%. The setting of G is a policy/performance issue that we investigate here.

Downstream/Upstream link (Mbps/Kbps)	Downstream/Upstream link Throughput (Kbps)	Downstream/Upstream link Link Utilization	K
8/64	18.8/62	0.235%/98%	5
8/256	49/252	0.612%/98%	1.25
8/384	126/376	1.57%/98%	0.5

Table 4: Bi-directional Single Flow Simulation without SmartAckDropper

Table 5 shows performance using a combination of CBQ, SAD and AR for different values of G and degrees of asymmetry. Aggregate (packet + ack) throughputs for both downstream and upstream directions are shown. Note that the performance in all cases has increased *by a factor of 10-20* compared to corresponding values in Table 4. Furthermore, as G increases from 0.1 to 0.9, the upstream aggregate throughput increases (and in some cases nearly doubles or triples) at the expense of downstream throughput (20-50% throughput loss), hence leading to a policy/performance tradeoff. In 8Mbps/256Kbps and 8Mbps/384Kbps (G=0.1 G=0.5) cases, the throughputs are almost equal is because upstream ack is saturated and the upstream data packet flow captures almost same bandwidth due to idle periods in the upstream ack stream.

Downstream/Upstream link(Mbps/Kbps)	Aggregate Throughput (bps/bps)	G
8/64	5.23M/22.4K	0.1
8/64	3.5M/32K	0.5
8/64	719K/56K	0.9
Downstream/Upstream link	Aggregate Throughput	G
8/256	6.61M/202K	0.1
8/256	6.61M/202K	0.5
8/256	2.97M/228K	0.9
Downstream/Upstream link	Aggregate Throughput	G
8/384	7.11M/326K	0.1
8/384	7.11M/326K	0.5
8/384	4.11M/342K	0.9

Table 5: Effect of SAD+AR+CBQ

#### 4.1.2 Loss and delay simulations

In most wireless network lossy links and delay links can cause throughput degradation due to bit errors and handoffs. In this section, we demonstrate that SAD results in significant performance improvements in these two environments.

##### • Downstream data loss simulation

In this simulation, we build a network loss module in downstream links, and set random loss level rate from  $10^{-4}$  to  $10^{-1}$ . SAD improves throughput a lot when the downstream loss rate is lower than  $10^{-2}$ . The throughput is down if loss rate is higher than  $10^{-2}$ . This is because SAD was not specially modified to handle interactions with packet loss and duplicated acks. Therefore, more spurious timeouts are generated.

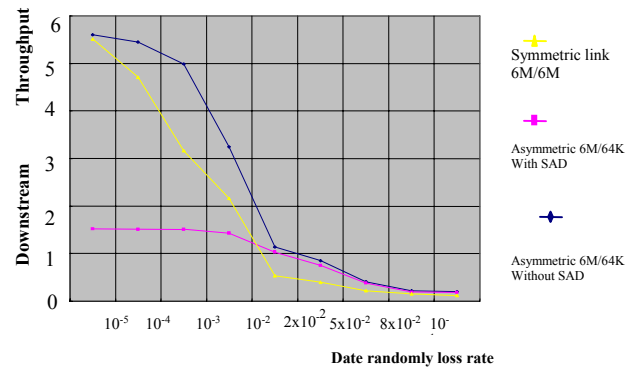


Chart 1: Uni-directional Single Flow Simulation on Loss Link

##### • Link delay simulation

We change link delay time in this simulation. Chart 2 shows that SAD still can improve performance on delay link ranging from 1ms to 100ms.

#### 4.2 VxWorks Implementation

This implementation is based on VxWorks: a real-time operating system running on Pentium platform. We first emulate asymmetric bandwidth and then implement SmartAckDropper on VxWorks testbed. Finally, we evaluate implementation results. The implementation source code is available in [9].




Chart 2: Uni-directional Single Flow Simulation on Delay Link

