Chapter 9

Results of Simulating TRUMP and RBCC in REAL

9.1 Introduction

Testing of the rate-based congestion control framework was performed on a number of hand-crafted network scenarios, as well as a large number of pseudo-randomly generated scenarios. The hand-crafted scenarios are described here.

The purpose of testing TRUMP/RBCC on hand-crafted scenarios is to demonstrate its behaviour where the points and severity of congestion can be determined in advance. As well, specific characteristics of the framework (such as fairness, ability to deal with route changes, ability to work with desired rates from sources) can be checked.

9.2 Description of Scenarios

Nine scenarios with different characteristics were crafted:

Scenario 1: A simple scenario with one traffic flow, one route and one congestion point.

Scenario 2: A scenario with several traffic flows and one congestion point. Flows start at different times, and the optimal flow rates at any time can be calculated.

Scenario 3: Several flows compete for a high-latency low-bandwidth link. This scenario should reveal any congestion problems caused by delay.

Scenario 4: Several flows come on-line at different times to compete for a low-bandwidth link. This scenario should reveal any temporal unfairness in rate allocation.

Scenario 5: Several flows with different route lengths compete for a low-bandwidth link. In this scenario, there are many cross-traffic flows with very short route lengths. This scenario should also reveal any spatial unfairness in rate allocation.

Scenario 6: A single long-haul traffic flow competes against a number of short-haul cross-traffic flows. Again, this scenario should also reveal any spatial unfairness in rate allocation.

Scenario 7: A scenario similar to Scenario 2, with packet sizes evenly distributed across a range. This should show any problems with varying packet sizes in TRUMP/RBCC.

Scenario 8: A scenario similar to Scenario 2, but with a route change during the scenario. This scenario should show if, and how well, TRUMP/RBCC can adapt to a route change in traffic flows.

Scenario 9: A modification to Scenario 6 to examine the effect of Rate Quenches on large selective acknowledgment delays.

9.2.1 Scenario Variations

There are three variations of each scenario:

TRUMP/RBCC All sources use TRUMP and all destinations use the TRUMP sink code. All routers use RBCC.

TCP Reno All sources use TCP Reno and all destinations use the standard REAL sink code. All routers are standard REAL routers.

TCP Vegas All sources use TCP Vegas and all destinations use the standard REAL sink code. All routers are standard REAL routers.

Nearly all scenarios use the same global parameters (the `real_params` sections shown in Appendix E). The most important parameters are given below:

- Acknowledgment packets are 40 octets in length, the same as TCP/IP acknowledgment packets in the Internet.
- Data packets are 1500 octets in length.
- The buffer size for the output interfaces in each router is 15,000 octets in length.
- For TCP Reno and TCP Vegas, the maximum window size is 600 segments. No source in the scenarios simulated reaches this value.
- TRUMP performs selective acknowledgments, with 1 data packet acknowledged in every acknowledgment packet. As will be seen in Chapter 10, >1 packets acknowledged per ack packet causes higher end-to-end and round-trip variance.
- RBCC scales an output interface's bandwidth by 0.75 when the output buffer for that interface reaches 5 packets or more. This congestion avoidance heuristic is described in Section 6.3.6.
- Routers use First-Come-First-Served as the packet selection method and Drop Tail (i.e the last packet to arrive) as the packet dropping method.
- Utilisation for each linked is averaged over a period of 2 seconds. The data obtained by REAL for link utilisation is coarse. REAL periodically calculates the link utilisation by dividing the amount of data transmitted over the link by the time period. If the period is too small, the utilisation information is inaccurate, being sensitive to packet scheduling etc. If the period is too long, the utilisation value is more accurate, but is the average utilisation over a larger time period.

Scenarios which do not use the parameters given above will be noted. The REAL input files containing these scenarios are given in Appendix E.

9.2.2 Tabled Results

In this chapter, a per-source table of results will be given for each scenario. The columns in the table hold:

- The end time of the source's data transmission, in seconds.
- The average data transmission rate, in bits per second. This includes the header in each packet, and not just the data payload.
- The average end-to-end time in seconds, and the standard deviation of the end-to-end time.
- The average round-trip time in seconds, and the standard deviation of the round-trip time.
- The number of data packets lost, and the number of data packets that were retransmitted.

9.3 Scenario 1

Scenario 1 is shown in Figure 16:

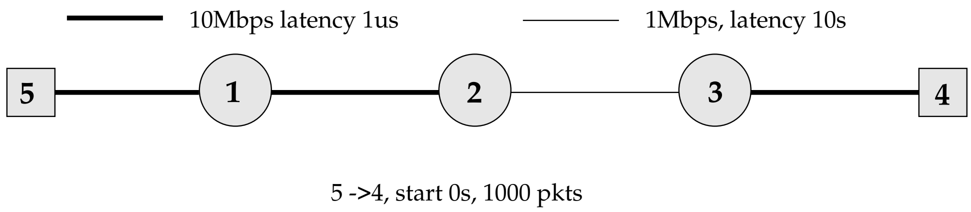

Figure 16: Scenario 1

This simple scenario is designed to show the transmission characteristics of the different transport protocols when there are no conflicting sources. The low-speed link 2 $\rightarrow$ 3 will limit the data rate of the data source at node 5. Router 2 may need to buffer packets if the source of traffic tries to exceed the available bandwidth.

Throughout this chapter, diagrams of network scenarios show squares as hosts, circles as routers, thick lines as high-speed links and thin lines as low-speed links. The speed and latency of each link is also given, as is the start time and number of data packets in each traffic flow.

9.3.1 Transmission Sequence Numbers

With all source types, node 5 begins transmission at 0 seconds. With TRUMP, the return handshake packet from node 4 reaches 5 at time 0.0109 seconds with a **Return Rate** value of exactly 1Mbps. Node 5 transmits data packets at this rate until time 12.0109, as shown in Figure 17.

TCP Reno does not perform a handshake, and begins data transmission at 0 seconds[1]. Acknowledgment packets open the window size (and hence the transmission rate) until a packet is lost by router 2 at time 0.154. The negative acknowledgment reaches the source

[1]In fact, for all scenarios, TCP Reno and TCP Vegas get the jump on TRUMP by transmitting without performing any handshaking. Similarly, neither Reno nor Vegas perform connection termination.

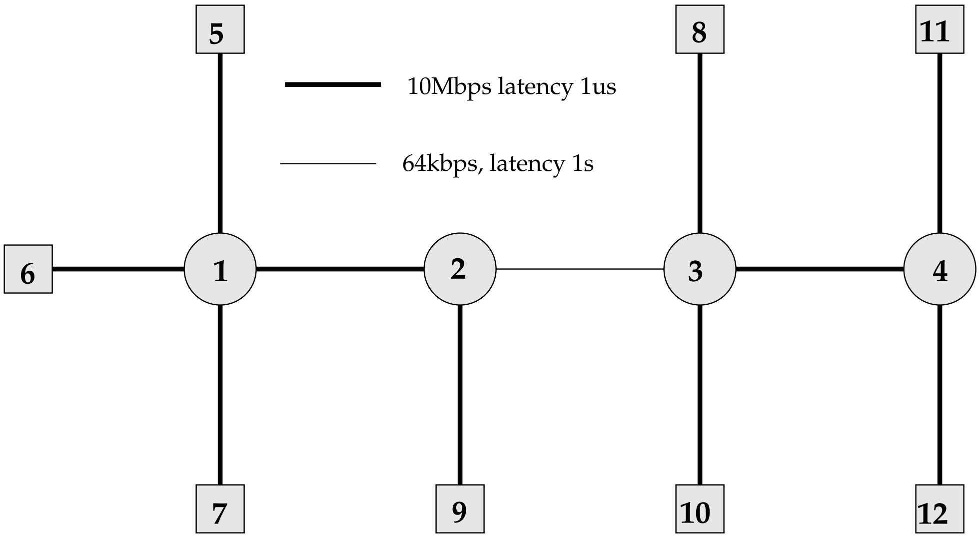

Figure 20: Scenario 2

9.4.1 Optimum Rates for Scenario 2

Given Scenario 2, it is simple to compute the optimum rates for each of the sources; the calculations are given in Appendix F. The following table summarises the rates; the parameter t is the simulation time in seconds. The value Off_x indicates the time that source x stops transmitting data. These optimum rates provide a baseline to compare the rates calculated by the RBCC scheme.

Time t (seconds)	**Optimum Rate (bps)**			
	Source 5	Source 7	Source 8	Source 6
$0 <= t < 12$	64,000			
$12 <= t < 18$	32,000	32,000		
$18 <= t <= 140$	31,168.8	31,168.8	62,337.7	
$140 <= t <= Off_6$	20,779.2	20,779.2	62,337.7	20,779.2
$Off_6 <= t <= Off_8$	31,168.8	31,168.8	62,337.7	
$Off_8 <= t$	32,000	32,000		

9.4.2 Transmission Sequence Numbers

Packets are transmitted by a source in monotonically increasing sequence number, barring retransmissions, with the derivative of sequence number versus time giving the effective transmission rate of a source. Figure 21 show the sequence number for each packet transmitted by sources 5 and 6 using TCP Reno, TCP Vegas and TRUMP.

Time t (seconds)	Optimum Rate (bps)			
	Source 5	Source 7	Source 8	Source 6
$0 <= t < 12$	64,000			
$12 <= t < 18$	32,000	32,000		
$18 <= t <= 140$	31,168.8	31,168.8	62,337.7	
$140 <= t <= Off_6$	22,168.8	22,168.8	62,337.7	18,000
$Off_6 <= t <= Off_8$	31,168.8	31,168.8	62,337.7	
$Off_8 <= t$	32,000	32,000		

Although flow $7 \rightarrow 12$ is bandwidth limited and flow $5 \rightarrow 10$ is not, they should be treated equally when $7 \rightarrow 12$ is allocated a rate of 32,000 bps or less.

The rates calculated by the distributed RBCC algorithms are:

Time t (seconds)	ABT Allocated Rate (bps)			
	Source 5	Source 7	Source 8	Source 6
2.02	64000.0			
14.15	32000.0	32000.0		
22.94	31180.0	31180.0	61440.0	
148.3	22180.0	22180.0	60960.0	18000.0
215.2	32000.0	32000.0	61440.0	

Again, the rates set by RBCC are very close to the optimum rates. As well, the bandwidth limited flows do not have their desired rates exceeded, and they are treated equally to the other flows when bandwidth must be shared.

9.6 Scenario 3

Scenario 3 is shown in Figure 29:

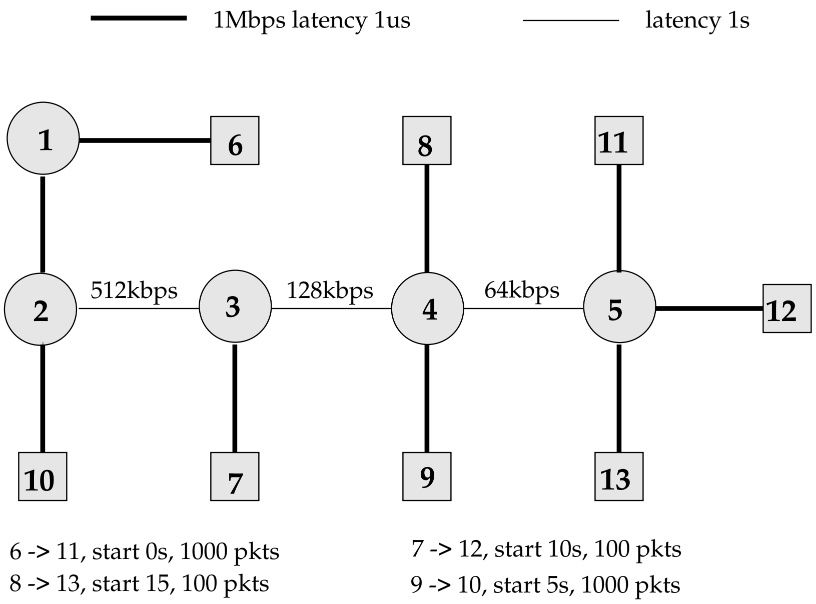

Figure 29: Scenario 3

This scenario is designed to see how well existing traffic sources react to other sources starting and stopping. The three low-speed links have very long latencies, and different capacities from each other. Three traffic flows are from left to right, and the $9 \rightarrow 10$ flow is from right to left. The start times for successive sources is close to the longest round-trip time.

9.6.1 Transmission Sequence Numbers

In Scenario 3, the TCP Reno sources 6 and 9 exhibit an exponentially increasing sequence number curve (and hence data rate curve). Both TRUMP's and Vegas' sequence number curves are linear. Note that Vegas' source 6 has not finished transmitting at time 2000 seconds. The per-source summary table below gives the end times and average data rates for all sources.

Source Node	End Time	Avg Rate	End to end Avg	End to end S.Dev	Round-trip Avg	Round-trip S.Dev	Pkts Lost	Pkts Retx'd
TRUMP 6	244.1	52077	3.41	0.31	6.44	0.33	0	33
Vegas 6	>2000	11919	3.41	0.14	2.10	0.27	0	1055
Reno 6	361.6	34811	4.07	0.56	7.04	0.51	3	5
TRUMP 7	72.7	22359	3.18	0.23	5.23	0.23	0	9
Vegas 7	179.1	14762	2.58	0.36	1.80	0.56	0	106
Reno 7	141.3	10324	3.25	0.47	5.15	0.49	4	6
TRUMP 8	75.5	21538	2.24	0.15	3.24	0.15	0	5
Vegas 8	53.5	35498	1.82	0.42	2.70	0.51	0	3
Reno 8	243.5	35558	1.98	0.49	2.60	0.39	21	27
TRUMP 9	112.9	120597	2.12	0.00	4.15	0.04	0	43
Vegas 9	1403.7	18351	2.17	0.05	1.32	0.28	0	1137
Reno 9	192.7	66890	2.15	0.05	4.16	0.10	0	1

TRUMP/RBCC and Vegas lose no packets, whereas Reno does. Note, however, that Vegas exhibits a large number of retransmissions: this scenario seems to reveal a defect in Vegas' retransmission timer strategy. End-to-end and round-trip deviation is generally lower for TRUMP/RBCC than for either TCP version. TRUMP sources have higher transmission rates than their equivalent TCP sources.

9.6.2 Router Queue Lengths

With TRUMP sources, no routers are congested[3] except for router 4, as shown in Figure 30. Its total buffer queue size for the $4 \rightarrow 5$ link is above 2 packets for time 17 to 76 seconds, spending most of the time at 6 packets queued.

[3] i.e. buffer sizes do not exceed one packet.

Source Node	End Time	Avg Rate	End to end Avg	End to end S.Dev	Round-trip Avg	Round-trip S.Dev	Pkts Lost	Pkts Retx'd
TRUMP 1	1497.6	8037	1.30	0.01	2.30	0.01	0	1
Vegas 1	1782.2	10126	2.69	0.33	3.28	0.80	288	496
Reno 1	1498.0	8962	2.43	0.56	3.33	0.26	106	112
TRUMP 2	1501.0	8024	1.42	0.01	2.42	0.01	0	1
Vegas 2	1736.5	10343	2.70	0.30	3.14	0.93	284	491
Reno 2	1529.5	8720	2.44	0.54	3.36	0.24	101	106
TRUMP 3	1503.1	8018	1.36	0.01	2.37	0.01	0	1
Vegas 3	1609.0	11604	2.70	0.31	3.03	0.98	278	550
Reno 3	1500.8	8941	2.42	0.58	3.31	0.23	104	110
TRUMP 4	1504.7	8015	1.29	0.01	2.29	0.01	0	1
Vegas 4	1756.0	9404	2.70	0.31	3.59	0.34	248	370
Reno 4	1458.5	9067	2.46	0.54	3.36	0.21	94	95
TRUMP 5	1506.2	8013	1.53	0.02	2.53	0.02	0	1
Vegas 5	1816.2	9461	2.62	0.44	3.40	0.62	269	419
Reno 5	1433.8	9213	2.44	0.57	3.32	0.20	94	95
TRUMP 6	1508.2	8016	1.40	0.02	2.40	0.02	0	2
Vegas 6	1800.7	9408	2.69	0.32	3.45	0.57	265	403
Reno 6	1575.6	8456	2.43	0.55	3.34	0.30	96	96
TRUMP 7	1508.9	8017	1.22	0.00	2.22	0.00	0	2
Vegas 7	1757.8	9904	2.71	0.29	3.51	0.48	278	439
Reno 7	1401.5	9397	2.44	0.56	3.34	0.17	90	90
TRUMP 8	1509.5	8019	1.56	0.03	2.56	0.03	0	2
Vegas 8	1796.0	11154	2.69	0.31	2.67	1.02	314	656
Reno 8	1584.2	8458	2.37	0.59	3.33	0.36	97	97

9.8 Scenario 5

Scenario 5 is shown in Figure 39. This simulation models a number of traffic flows, some of which flow across different sections of the 1M, 100k, 64k, 1M low-speed link. A congestion control scheme should not penalise one traffic flow for passing through more intermediary routers than another traffic flow.

Source Node	End Time	Avg Rate	End to end Avg	End to end S.Dev	Round-trip Avg	Round-trip S.Dev	Pkts Lost	Pkts Retx'd
TRUMP 11	177.8	67973	0.19	0.04	0.25	0.08	0	1
Vegas 11	170.5	71433	0.60	0.18	0.69	0.14	0	5
Reno 11	156.3	81758	0.81	0.27	0.87	0.18	25	50
TRUMP 12	56.4	37448	1.15	0.08	2.17	0.11	0	4
Vegas 12	127.4	21318	2.25	0.50	1.98	1.34	5	85
Reno 12	147.3	11398	2.31	0.54	3.53	0.46	21	17
TRUMP 13	441.0	30220	0.99	0.01	2.02	0.12	0	5
Vegas 13	976.0	27527	1.55	0.46	1.02	0.45	27	1141
Reno 13	465.9	30150	1.66	0.51	3.10	0.37	31	60
TRUMP 14	42.2	1008221	0.02	0.00	0.04	0.00	0	2
Vegas 14	42.3	1003458	0.06	0.01	0.07	0.01	0	0
Reno 14	49.6	934329	0.10	0.01	0.05	0.02	0	24
TRUMP 15	448.2	30539	1.18	0.04	2.39	0.15	0	11
Vegas 15	1105.2	24529	1.75	0.50	1.20	0.33	32	1152
Reno 15	633.9	22035	1.87	0.58	3.45	0.39	73	67

9.9 Scenario 6

Scenario 6 is shown in Figure 44:

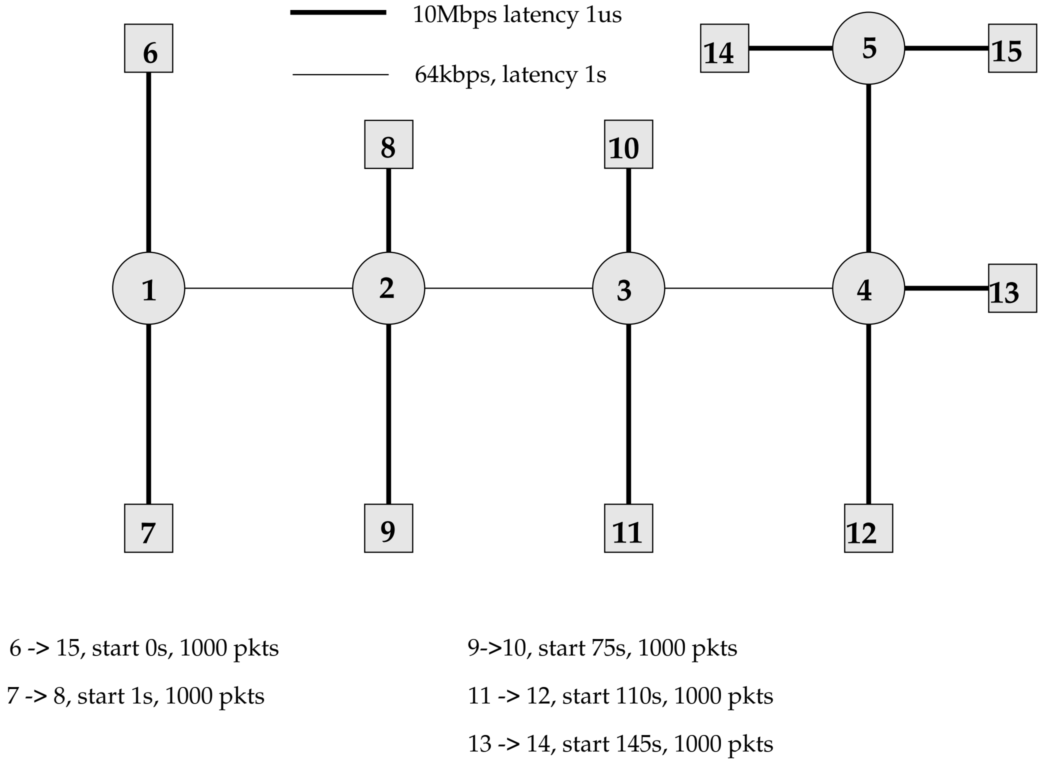

Figure 44: Scenario 6

This scenario models a long-haul traffic flow, 6 → 15, which must compete against a number of short-haul cross-traffic flows. The end times for all sources are given below.

Flow	Range
5 → 10	[70:1500]
6 → 11	[400:900]
7 → 12	[700:1500]
8 → 9	[900:1500]

Unfortunately, the version of TCP Reno and TCP Vegas that are available in REAL 4.0 do not have the ability to alter their packet sizes on the fly, so no comparisons can be drawn with TRUMP/RBCC and TCP in this scenario.

The usual per-source results for Scenario 7 are given in the following table.

Source Node	End Time	Avg Rate	End to end		Round-trip		Pkts	
			Avg	S.Dev	Avg	S.Dev	Lost	Retx'd
TRUMP 5	207.2	58939	1.16	0.07	2.20	0.09	0	8
TRUMP 6	168.2	49679	1.19	0.07	2.22	0.08	0	8
TRUMP 7	260.1	49606	1.20	0.06	2.26	0.08	0	17
TRUMP 8	274.4	47842	1.15	0.02	2.28	0.08	0	13

Packet loss is nonexistent, and end-to-end and round-trip deviation is low. Packet queue lengths in the bottleneck router is higher than for Scenario 2, and fluctuates more, as shown in Figure 46. This is a consequence of the sizes and earlier arrival times for the smaller packets. This scenario shows that a congestion avoidance mechanism which relies simply on the number of packets queued will sometimes incorrectly perceive congestion.

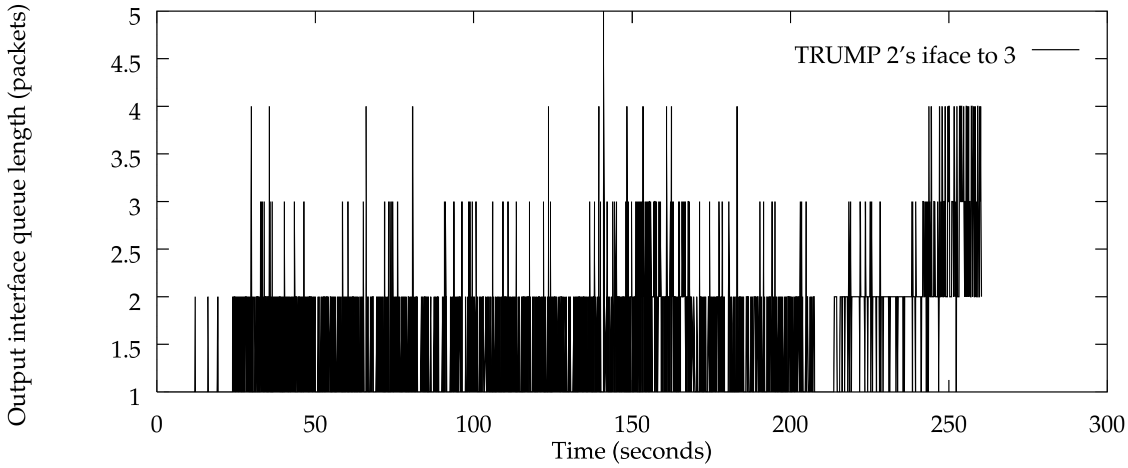

Figure 46: TRUMP Queue Lengths in Scenario 7

The effect of the small packets and the larger router queue sizes has little effect on the transmission rates passed back by RBCC to the traffic sources, as shown in Figure 47. The rate are, as usual, stable, but with a number of dips where the RBCC congestion avoidance mechanism has fired. Link utilisation is essentially unchanged from Scenario 2.

Packet Losses and Retransmissions Overall, TRUMP loses *no* packets in any of the scenarios. Vegas loses many fewer packets than Reno, except for Scenario 4 where Vegas loses over 3 times as many packets as Reno. However, Vegas and Reno loses 2,956 and 1,563 packets, respectively. On this basis alone, TRUMP/RBCC provides better congestion control than TCP.

The amount of packet retransmission varies greatly between scenarios, but in general TRUMP's level of packet retransmission is roughly of the same magnitude as TCP Reno's. Vegas, on the other hand, appears to aggressively retransmit packets, which indicates that its timeout strategy is still incorrect. Note that TRUMP does not retransmit packets because of packet losses: instead, it lazily retransmits packets until their acknowledgments are received. This is done to avoid setting a round-trip timer when there is no new data to transmit.

Router Queue Lengths TRUMP/RBCC generally keeps router queue lengths down around 1 or 2 packets, except where new traffic flows came on line: the delay to return new sustainable rates to all traffic flows allows router queue lengths to grow. The primitive 'threshold' congestion avoidance scheme in RBCC is adequate to bring the queue levels back down to 1 or 2 packets. In contrast, both TCP versions keep queue lengths in bottleneck routers above 2 packets most of the time. This adds to the delay of packet transmission, and makes each router more susceptible to short-term congestion which fills the remaining buffer space and causes lost packets. On the basis of router queue lengths, TRUMP/RBCC appears to give good congestion control.

End-to-end & Round-trip Times The higher router queue lengths in TCP increases the average end-to-end and round-trip times for traffic flows. Due to the lower router queue lengths in TRUMP/RBCC, average end-to-end and round-trip times are also lower. The 'bursty' window size increases made by both TCP flavours increases router queue lengths in the short-term, which adds variance to end-to-end and round-trip times. Ironically, this in turn makes TCP's round-trip time estimation more difficult. Not only does TRUMP/RBCC not require a round-trip timer, but its Leaky Bucket packet admission works to reduce variance in end-to-end and round-trip times.

Although this feature is not directly related to congestion control, it is a boon to applications such as voice and video transmission where low end-to-end time variance is desired.

Link Utilisation and RBCC Rate Allocations RBCC appears to derive near-optimal transmission rates for all traffic sources. This keeps 'bottleneck' links operating at near 100% utilisation. TCP's Slow Start and bandwidth-probing strategies either keeps bottleneck links operating either below their capacity, or offers too much load for links to transmit. The result is either poor utilisation or packet loss. As well, for some scenarios, TCP did not probe and find extra bandwidth when it became available. The reason for this behaviour is unknown.

Scenario Variations Finally, the proposed congestion control framework, with TRUMP and RBCC as component functions, was able to cope with finite desired rates from traffic sources, ranges of packet sizes from sources, and low-frequency route changes.