Remote Hardware Control:
Mission-critical Application Performance
with and without Traffic Differentiation

DRAFT Editor: Tiziana Ferrari
Italian National Institute for Nuclear Physics (INFN-CNAF)

QUAdiS Project

1. Introduction

In this paper we analyze the sensitivity of high priority traffic to TCP burstiness and traffic load in different scenarios. The end-to-end performance of the high priority traffic – expressed by the *Round Trip Time* metric – is evaluated both in a wide-area test network connecting two test sites at INFN and FNAL and in a production environment.

The performance gain introduced by QoS features like classification, marking, policing and scheduling is quantified. In particular, two scheduling algorithms – Priority Queuing and Weighted Fair Queuing – are compared. For all the experimental results a specific application for hardware remote control in high energy physics – *Robin* [1] – is used.

In Section 2 we illustrate the network test bed deployed for traffic differentiation testing coupled with the hardware remote monitoring application. In particular, Section 3 presents the end-to-end performance estimated under different traffic scenarios without traffic differentiation, while Section 4 reports on the different end-to-end performance experienced when different queuing techniques are enabled, namely Priority Queuing and Weighted Fair Queuing. In Section 5 we analyze the performance of the same application when running on a standard production data-path and we comment on several properties of over-provisioned congestion-free production networks.

2. Network Layout and Baseline Measurements

The test bed is composed of three routers connected in a chain by point-to-point ATM constant bit rate (CBR) connections of 2 Mbps (the link capacity includes the ATM protocol overhead) – see Fig. 1. Only Router 1[1] and Router 3 are congestion points, since traffic is injected by workstations directly connected to the above-mentioned devices. On the other hand, Router 2 is "transparent", i.e. we assume that the queuing delay introduced by queuing in Router2 is infinitely small in comparison with the overall queuing delay introduced by the output interfaces of Router 1 and 2.

The VME controller is directly connected to Router 3, while the Robin client is directly connected to Router 1. Traffic exchanged between the Robin client and the VME controller is packet-loss and delay-sensitive, as such it is subject to preferential treatment over background traffic exchanged by other hosts in the test bed. In the following we will refer to high priority traffic with term *Expedited Forwarding* traffic (EF traffic) in

---

[1] Router 1 runs IOS experimental version 12.0(6.5)T7.

compliance with diffserv terminology. Expedited Forwarding is the Per Hop Behavior which is more suitable for delay and jitter-sensitive traffic.

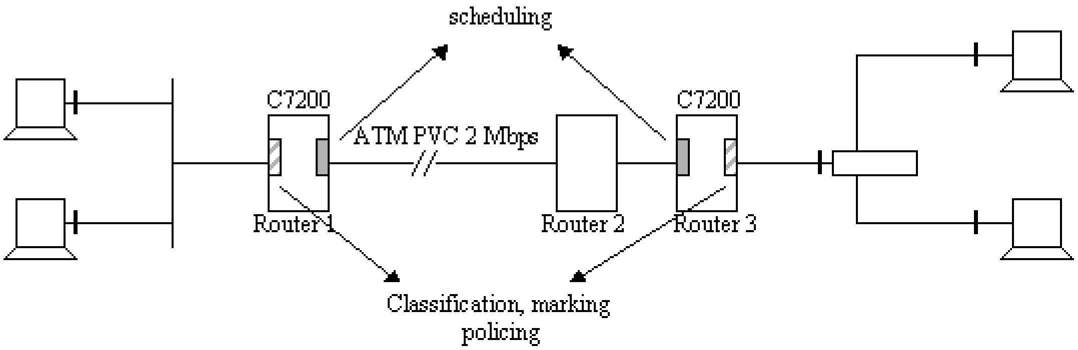

Figure 1: Network test bed

Round Trip Time (RTT) is the metric used for end-to-end performance analysis. RTT of best-effort 64-by long packets varies in the range [124.3, 127.0] msec, while the average RTT is 124.5 msec. The estimated loss in presence of full line rate UDP traffic is null[2].

The TCP flow generated by the test application is based on a maximum TCP socket size of 32696 by, while the average socket size is 4096 by[3].

3. Best-Effort Traffic Load

In order to quantify the benefits of traffic differentiation in terms of end-to-end delay we measured RTT in a completely best-effort scenario by increasing gradually the background traffic load. Since delay is highly sensible to instantaneous packet burst accumulation, we have run two sets of tests: with UDP and TCP background traffic. Test results show that the end-to-end performance in the two cases varies greatly.

UDP Background Traffic

Experimental results show that RTT is a function of UDP best-effort background load as shown in Fig.2. We have injected a variable number of one-way independent well-shaped CBR UDP streams, so that the overall injected UDP rate varies in the range [500, 1000, 1500, 2000] Kbps (the rates are ATM overhead-inclusive).

---

[2] The UDP traffic the measurement refers to was consists of a CBR stream of 210 pack/sec, with payload packet size equal to 1000 by. The overall throughput (including the ATM overhead) is 1.957 Mbps.
Unlike UDP, in case of long-lived TCP connections an asymmetric behavior was observed: TCP long-lived connections from INFN (Router 1) to FNAL (Router 2) could achieve a maximum throughput of approximately 600 Kbps – even in presence of multiple streams – while in the opposite direction the link could be completely saturated. The problem was probably due to a policing configuration error. Nevertheless, short lived TCP connections like the ones produced by our test application where not affected by the problem.

[3] The small average socket size is due to the fact that Robin generates short-lived TCP connections, i.e. TCP connections are not long enough to let the transport protocol increase the window size.

Fig. 2 shows that the RTT frequency distribution gets spread over a larger and larger range when the amount of competing traffic gradually increases.

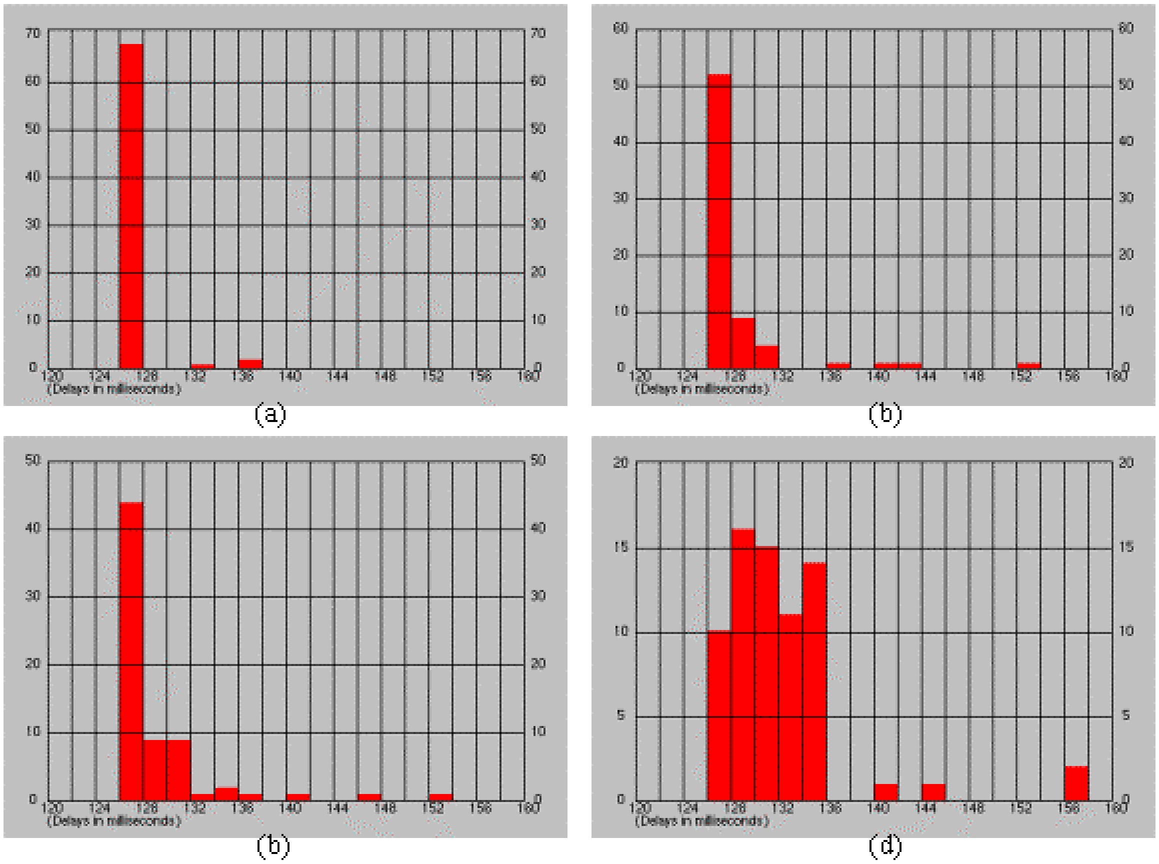

(a) (b)

(b) (d)

Figure 2: RTT in presence of 500 kbps (a), 1000 Kbps (b), 1500 kbps (c) and 2000 Kbps (d) of background UDP traffic – the horizontal axis display the number of RTT samples falling in a given RTT range –

TCP Background Traffic

The testing of the correlation between RTT and TCP is important to evaluate the impact of TCP burstiness even in absence of long-term congestion. We have augmented the burstiness of the background traffic by increasing the number of TCP streams from 4 to 32. The overall TCP average rate never exceeded 642 Kbps, i.e. 32% of the line rate. Streams were sourced by INFN and received by a test workstation at FNAL. The maximum send and receive TCP socket buffer size was 240,000 by.

We have analyzed the RTT frequency distribution of a population of approximately 100 samples. As Fig. 3 shows, even in absence of long term congestion, RTT is greatly affected by the number of TCP background concurrent streams, in fact the greater is the number of streams the more frequent is the presence of a TCP packet burst in the output queue of the ingress router. Results in Fig. 3 show that the Robin application is more sensitive to TCP background traffic rather than UDP traffic, since in the former case with 32 connections RTT can be up to 1.1 sec.

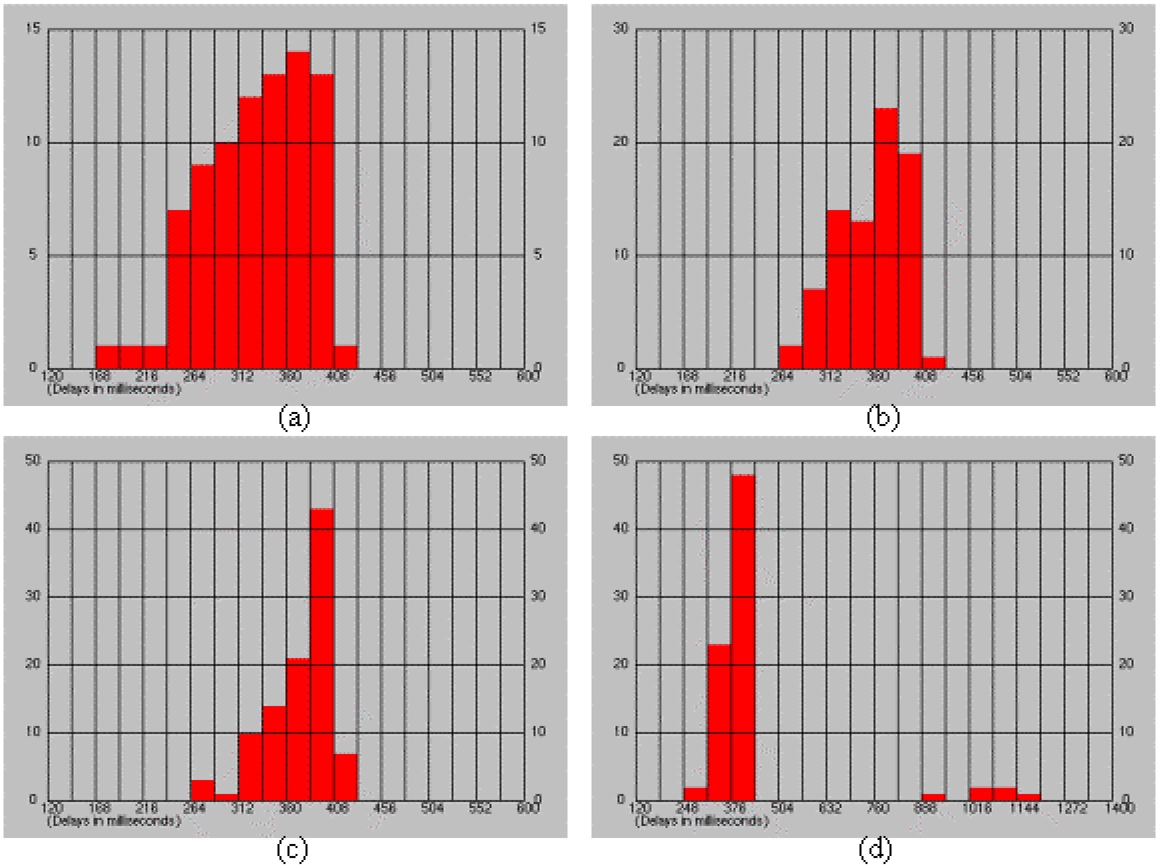

Figure 3: RTT of Robin packets in presence of TCP background streams: 4 TCP streams (a), 8 streams (b), 16 streams (c) and 32 streams (d)

RTT depends on the number of aggregation points on the data path: the longer the path, the greater is the probability that a reference packets gets queued at least once from source to receiver behind a TCP background stream. In addition, since TCP traffic is bi-directional, the overall performance of a Robin session also depends on the acknowledgement delay, which in its turn is affected by instantaneous or permanent congestion on the return path.

This latter consideration is supported by our experimental results: As shown by Fig. 4, f TCP background traffic is injected in both directions (from INFN to CNAF and vice versa[4]), then RTT is not only distributed over a larger range, but I addition some samples can experience a considerable delay up to 3.0 sec.

---

[4] The background TCP average load of traffic from FNAL to CNAF is equal to 2.013 Mbps.

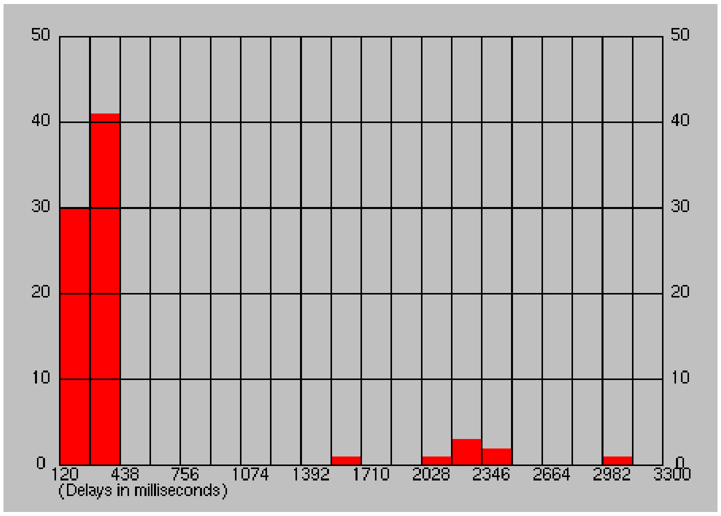

Figure 4: RTT in case of two-way TCP traffic (32 TCP streams from INFN to FNAL plus additional 16 streams from FNAL to INFN)

4. RTT Performance with Priority Queuing and WFQ

Given the baseline measurements presented in the previous paragraphs, we have introduced QoS Differentiated Services features (namely classification, marking, policing and scheduling[5]) to quantify the RTT performance gain introduced by QoS.

We have compared two different queuing mechanisms: strict Priority Queuing (PQ) and Weighted Fair Queuing (WFQ).

A mixture of TCP and UDP streams from INFN to FNAL was used, namely 6 UDP streams for an overall rate exceeding the line capacity – to produce permanent congestion so that scheduling is active for the whole duration of the test – plus additional 6 TCP streams. UDP and TCP flows are sourced by two different senders directly connected to Router 1, conversely, traffic from FNAL to INFN is produced by a set of 6 TCP streams[6].

In both cases a set of 8 different queues is used – each of them corresponding to a precedence value – and a minimum rate of 800 Kbps is assigned to *priority* traffic (i.e. Robin packets). With Priority Queuing just one Priority Queue is configured to store precedence 7 traffic, while the remaining queues are fed with packets with precedence in the range [0, 6]. Best-effort packets are assigned to one dedicated queue corresponding to precedence 0. On the other hand, with WFQ also Robin packets are placed in a WFQ queue, which is assigned a minimum bandwidth of 800 Kbps.

In both cases all queues in the scheduling module are constantly fed by traffic so that the edge router is a point of long-term congestion.

---

[5] Refer to Appendix B and C for details.

[6] Documentation on stream profiles is available in Appendix A.

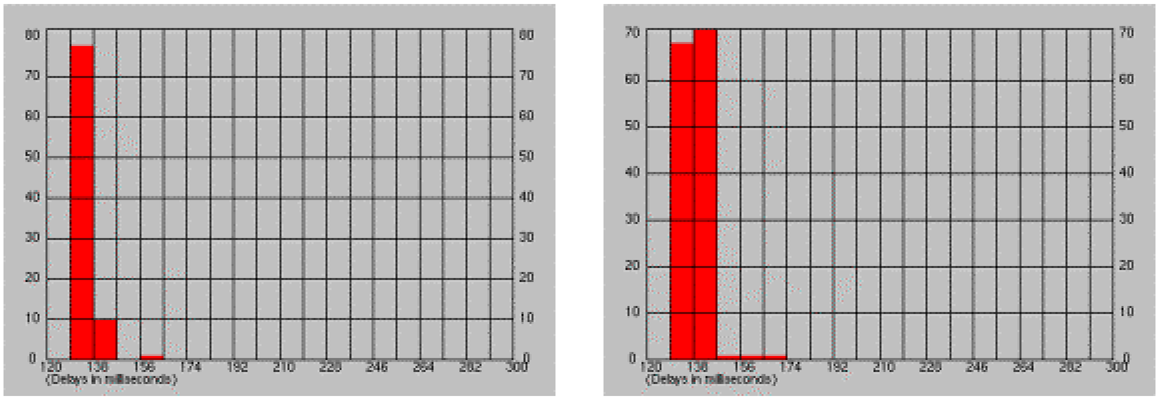

(a) (b)

Figure 5: RTT performance in case of PQ (a) and WFQ (b) with one-way traffic

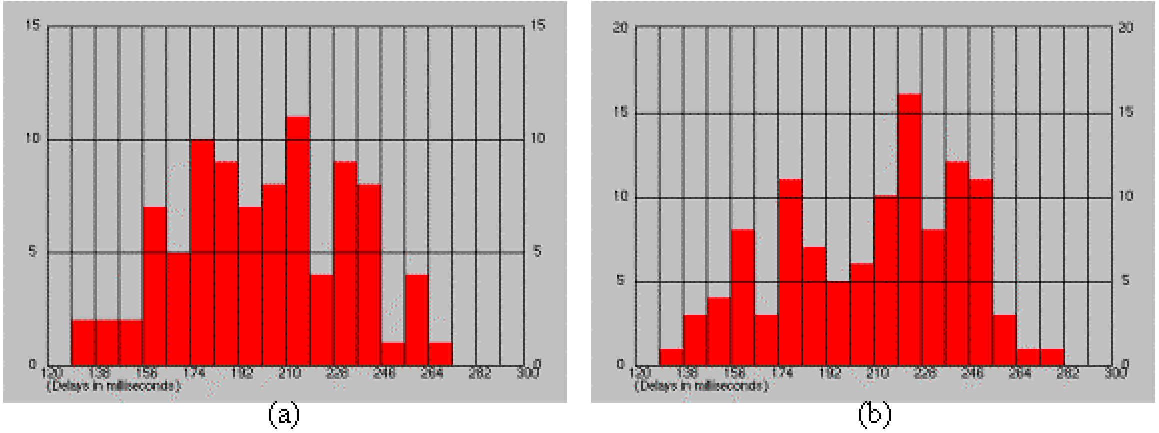

(a) (b)

Figure 6: RTT performance in case of PQ (a) and WFQ (b) with two-way traffic

As Fig.5 and Fig.6 show, with both one-way and two-way traffic RTT is minimized by Priority Queuing: In both cases while the minimum range value is the same, the maximum RTT value experienced is greater with WFQ. Note that in this test while priority traffic is fed into a WFQ queue, proper bandwidth over-provisioning is adopted to make sure that the Robin queue is drained at a proper rate and no queue build-up occurs (800 Kbps is well above the input rate).

This experimental result validates both the experimental results achieved in the TF-TANT framework [26] and the theoretical performance analysis presented in literature.

5. RTT in a Production Environment

For comparison's sake the performance of the remote monitoring application was estimated when run in a production environment. In this case our reference packets have to compete with production traffic for transmission. Bandwidth on the data path from INFN to FNAL – which goes through the GARR-B [2], TEN-155 [3,4] and ESnet networks – is largely over-provisioned, as indicatively shown in Fig. 7 by statistics of research traffic from GARR-B (PoP in Milano) to the New York PoP of TEN-155. On this data path peak traffic utilization is approximately 50% of the line capacity. However, RTT greatly depends on the time of the day; the average RTT is around 139.2 msec (see Appendix D).

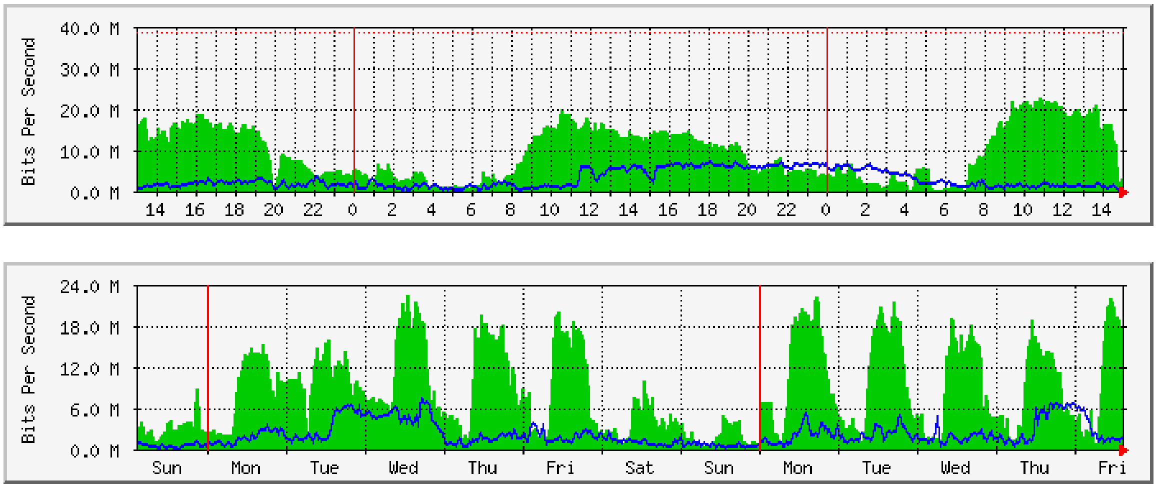

Figure 7: Daily (top) and weekly (bottom) research traffic statistics between GARR-B (Milano PoP) and the TEN-155 New York PoP[7]

No quality of service features were enabled anywhere on the path from the source to the destination, in other terms, no differentiation between traffic classes was deployed and reference packets received a best-effort treatment. No background traffic was injected, instead production traffic was used in its place.

We have precisely estimated the interaction between production traffic and remote control traffic in terms of RTT frequency distributions over a population of approximately 150 samples, i.e. we used sets of 150 client-server transactions for RTT analysis. Since the automatic execution of client-server transactions was not possible we had to restrict the measurement period to short intervals of tens of minutes and systematic monitoring of performance on the line was not possible. This implies that the probability of hitting a short-term congestion period was smaller.

We used the monitoring tool *tcptrace* [5] for an estimation of packet loss. Test results show that *no packet loss* is normally experienced by a set of subsequent 150 client-server transactions. However, while RTT is normally stable, a small percentage of transactions (2-5%) can still experience considerable delay depending on the time of the day, as shown by the RTT distribution tail in Fig. 8a.

Distribution in Fig. 8b is similar to distribution in Fig. 2b, which is computed when reference traffic packets are treated as best-effort with 50% of the line capacity occupied by background traffic. Conversely, it shows a longer tail than the RTT distribution in Fig. 5a, which was achieved when priority packets are served by a priority queue. This suggests that dynamic short-term congestion does occur even during very

---

[7] Traffic statistics are from the GARR-B monitoring site http://www.noc.garr.it/mrtg-II/

short monitoring intervals and that it can affect the application end-to-end performance.

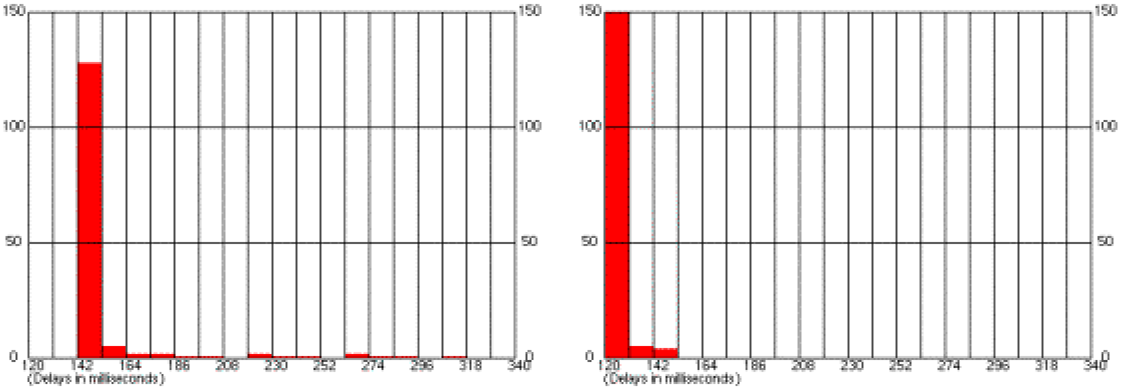

(a) Case1: 11 am (b) Case1: 3 pm

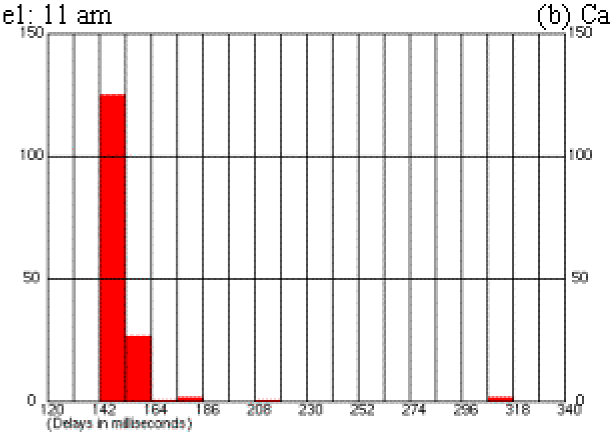

(c) Case3: 10 pm

Figure 8: RTT frequency distribution (expressed in number of packets) of remote control packets when transmitted on the production data path at different times of the day

6. Conclusions and future work

Experimental results show that Robin TCP sessions are both sensitive to traffic load and burstiness. Traffic burstiness is the major parameter since in absence of permanent line congestion, even for a long-term line capacity utilization as low as 32%, about 50% of high priority packets can experience a RTT which is more than twice the maximum RTT experienced in case of constant bit rate full line-capacity load.

Results also show that two-way differentiation is necessary for full high priority traffic protection and that priority queuing is the most suitable scheduling algorithm to minimize RTT, this validates the simulation and measurement-based analysis discussed in many analytical studies.

Another important result is that the RTT experienced on long–term congestion-free production networks can be less stable than RTT when priority queuing is enabled to differentiate traffic congestion periods. In particular, the end-to-end high priority traffic profile is comparable to the one achieved with a *constant bit rate* traffic load equal to 50% of the line capacity, while it behaves much better than the RTT measured with a small number of background *TCP* (bursty) connections in a controlled test network

TCP performance in the production environment and in the test bed can differ and have a different impact on end-to-end performance. In the former case, core routers –

which aggregate a huge amount of TCP streams – do not experience the TCP burstiness seen by edge devices – like in the testbed – and this contributes to minimize the probability of short-term congestion. Conversely, the input TCP burstiness is smoothed out at the output interface of an edge router at the cost of greater end-to-end delay – caused by queue congestion – and of packet loss. We can conclude, that traffic differentiation mechanisms are fundamental to preserve high priority traffic in each potential congestion point along a given data path.

Results suggest that edge routers are the devices more affected by traffic burstiness; however, the detailed measurement of packet loss percentage and of short-term congestion probability in an over-provisioned backbone infrastructure is required and is subject of future research.

7. Bibliography

[1] *ROBIN: The Rpc and Object Broker Interface*: http://www-b0.fnal.gov:8000/ROBIN/

[2] *GARR, The Italian Academic and Research Network*: http://www.garr.net/

[3] *The TEN-155 European Research Network*: http://www.dante.net/ten-155/

[4] *Quantum Project (the Quality network Technology for User-Oriented Multi-Media)*;http://www.dante.net/quantum/

[5] *Tcptrace Home Page (Ohio University):* http://jarok.cs.ohiou.edu/software/tcptrace/tcptrace.html

[6] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., Weiss, Z. W.: *An Architecture for Differentiated Services*, RFC 2475:

[7] Jacobson, V., Nichols, K., Poduri, K.: *An Expedited Forwarding PHB*, RFC 2598

[8] Nichols, K., Jacobson, V., Zhang, L.: *A Two-bit Differentiated Services Architecture for the Internet*.

[9] Nichols, K., Lake, S., Baker, F., Black, D.: *Definition of the Differentiated Services Field (DS Field) in the Ipv4 and Ipv6 Headers*; RFC 2474

[10] Bernet, Y., Smith, A., Blake, S.: *A Conceptual Model for Diffserv Routers*, diffserv draft, work in progress.

[11] *Differentiated Service*, http://www.ietf.org/html.charters/diffserv-charter.html

[12] *IP Performance Metrics*; http://www.ietf.org/html.charters/ippm-charter.html

[13] Almes, G., Kalidindi, S., Zekauskas, M.: *One-way Delay Metric for IPPM*, RFC 2679

[14] Demichelis, C., Chimento, P.: *Instantaneous Packet Delay Variation Metric for IPPM*, ippm draft, work in progress.

[15] Bradner, S.: *Benchmarking Terminology for Network Interconnection Devices*, RFC 1242

[16] Mills, D. L.: *Network Time Protocol (Version 3), Specification, Implementation and Analysis*, RFC 1305

[17] Guerin, R., Li, L., Nadas, S., Pan, P., Peris, V.: *The Cost of QoS Support in Edge Devices: An Experimental Study;* INFOCOMM 1999

[18] *The Multi-Generator (MGEN) Toolset*, Naval Research Laboratory (NRL) http://manimac.itd.nrl.navy.mil/MGEN/

[19] Prue, W., Postel, J.: *A Queuing Algorithm to Provide Type-of-Service for IP Links*, RFC 1046

[20] Zhang, H.: *Service Disciplines For Guaranteed Performance Service in Packet-Switching Networks*

[21] *The bandwidth guaranteed prioritized queuing and its implementations law*, K.L.E. Global Telecommunications Conference, 1997. GLOBECOM '97., IEEE Volume: 3 , 1997, Page(s) 1445-1449 vol3.

[22] Rönngren, R., Ayani, R.: *A comparative study of parallel and sequential priority queue algorithms*, ACM trans. Model. Comput. Simul. 7,2 (Apr. 1997), pages 157 – 209

[23] Parekh, A., Gallager, R.: *A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single-Node Case*; IEEE/ACM Transactions on Networking, Vol 1, No 3, June 1993

[24] Floyd, S, Jacobson, V.: *Link-sharing and Resource Management Models for Packet Networks*; ACM Transactions on Networking, Vol 3 No. 4, Aug 1995

[25] Ferrari, T. (Editor): *Differentiated Service Experiment Report*, TF-TANT interim report Jun 99 – Sep 99. http://www.cnaf.infn.it/~ferrari/tfng/ds/del-rep1.doc

[26] *The Joint DANTE/TERENA Task Force TF-TANT*; http://www.dante.net/tf-tant/

APPENDIX A: TCP and UDP traffic profile for Robin tests *with* QoS

TRAFFIC PROFILE from INFN to FNAL:

6 UDP streams (each marked with a different precedence value):

```
mgen -b 192.190.218.2:50001 -p 50001 -d 60000 -i ge0:6 -s 500 -r 50  &
mgen -b 192.190.218.2:50002 -p 50002 -d 60000 -i ge0:6 -s 500 -r 50  &
mgen -b 192.190.218.2:50003 -p 50003 -d 60000 -i ge0:6 -s 500 -r 50  &
mgen -b 192.190.218.2:50004 -p 50004 -d 60000 -i ge0:6 -s 500 -r 50  &
mgen -b 192.190.218.2:50005 -p 50005 -d 60000 -i ge0:6 -s 500 -r 50  &
mgen -b 192.190.218.2:50006 -p 50006 -d 60000 -i ge0:6 -s 500 -r 50  &
mgen -b 192.190.218.2:50011 -p 50011 -d 60000 -i ge0:6 -s 500 -r 200 &
```

6 TCP streams (each marked with precedence 6):

```
netperf -H 192.190.218.2 -l 3600 -p 40000 -- -s240000,240000 -S240000,240000 &
netperf -H 192.190.218.2 -l 3600 -p 40001 -- -s240000,240000 -S240000,240000 &
netperf -H 192.190.218.2 -l 3600 -p 40002 -- -s240000,240000 -S240000,240000 &
netperf -H 192.190.218.2 -l 3600 -p 40003 -- -s240000,240000 -S240000,240000 &
netperf -H 192.190.218.2 -l 3600 -p 40004 -- -s240000,240000 -S240000,240000 &
netperf -H 192.190.218.2 -l 3600 -p 40005 -- -s240000,240000 -S240000,240000 &
```

TRAFFIC PROFILE from FNAL to INFN (if present):

```
netperf -H 192.168.190.12 -l 60000 -p 40000 -- -s240000,240000 -S240000,240000 &
netperf -H 192.168.190.12 -l 60000 -p 40001 -- -s240000,240000 -S240000,240000 &
netperf -H 192.168.190.12 -l 60000 -p 40002 -- -s240000,240000 -S240000,240000 &
netperf -H 192.168.190.12 -l 60000 -p 40003 -- -s240000,240000 -S240000,240000 &
netperf -H 192.168.190.12 -l 60000 -p 40004 -- -s240000,240000 -S240000,240000 &
netperf -H 192.168.190.11 -l 60000 -p 40001 -- -s240000,240000 -S240000,240000 &
```

APPENDIX B: WFQ router configuration

```
policy-map cdf-wfq   /* scheduling module, WFQ */
  class pre7         /* Robin – high priority -  WFQ queue */
   bandwidth 800
  class pre6
   bandwidth 100
   queue-limit 10
  class pre5
   bandwidth 100
   queue-limit 10
  class pre4
   bandwidth 100
   queue-limit 10
  class pre3
   bandwidth 100
   queue-limit 10
  class pre2
   bandwidth 100
   queue-limit 10
  class pre1
   bandwidth 100
   queue-limit 10
  class class-default
```

```
   bandwidth 100

class-map pre7             /* precedence 7 traffic */
  match access-group 180
class-map pre6             /* precedence 6 traffic */
  match access-group 181
class-map pre5             /* precedence 5 traffic */
  match access-group 105
class-map pre4             /* precedence 4 traffic */
  match access-group 183
class-map pre3             /* precedence 3 traffic */
  match access-group 184
class-map pre2             /* precedence 2 traffic */
  match access-group 185
class-map pre1             /* precedence 1 traffic */
  match access-group 186

interface FastEthernet0/0
 description test LAN
 ip address 192.168.72.1 255.255.255.0 secondary
 ip address 192.168.174.1 255.255.255.0 secondary
 ip address 192.168.184.1 255.255.255.0 secondary
 ip address 192.65.183.129 255.255.255.240 secondary
 ip address 192.168.186.1 255.255.255.0 secondary
 ip address 131.154.98.253 255.255.255.0 secondary
 ip address 192.168.190.14 255.255.255.248 secondary
 ip address 192.168.73.1 255.255.255.0
 no ip directed-broadcast
/* classification, marking and policing */
 rate-limit input access-group 150 1000000 8000 8000 conform-action set-prec-transmit 7
exceed-action drop   /* Robin traffic */
 rate-limit input access-group 161 1000000 8000 8000 conform-action set-prec-transmit 1
exceed-action set-prec-transmit 1
 rate-limit input access-group 162 1000000 8000 8000 conform-action set-prec-transmit 2
exceed-action set-prec-transmit 2
 rate-limit input access-group 163 1000000 8000 8000 conform-action set-prec-transmit 3
exceed-action set-prec-transmit 3
 rate-limit input access-group 164 1000000 8000 8000 conform-action set-prec-transmit 4
exceed-action set-prec-transmit 4
 rate-limit input access-group 165 1000000 8000 8000 conform-action set-prec-transmit 5
exceed-action set-prec-transmit 5
 rate-limit input access-group 166 1000000 8000 8000 conform-action set-prec-transmit 6
exceed-action set-prec-transmit 6
 ip route-cache policy
 load-interval 30
 full-duplex

interface ATM1/0.108 point-to-point
 description to ESnet (CDF QoS testing)
 ip address 192.168.190.2 255.255.255.252
 no ip directed-broadcast
 pvc 4/108
  tx-ring-limit 5
  service-policy output cdf-wfq   /* scheduling is activated */
  vbr-nrt 2000 2000 1
  encapsulation aal5snap

access-list 150 permit ip any host 192.190.218.1
access-list 161 permit udp any host 192.190.218.2 eq 50001
access-list 162 permit udp any host 192.190.218.2 eq 50002
access-list 163 permit udp any host 192.190.218.2 eq 50003
access-list 164 permit udp any host 192.190.218.2 eq 50004
access-list 165 permit udp any host 192.190.218.2 eq 50005
access-list 166 permit tcp any host 192.190.218.2

access-list 180 permit ip any any precedence network
access-list 181 permit ip any any precedence internet
access-list 182 permit ip any any precedence critical
access-list 183 permit ip any any precedence flash-override
access-list 184 permit ip any any precedence flash
access-list 185 permit ip any any precedence immediate
```

```
access-list 186 permit ip any any precedence priority
```

APPENDIX C: PQ router configuration

```
class-map pre7              /* precedence 7 traffic */
  match access-group 180
class-map pre6              /* precedence 6 traffic */
  match access-group 181
class-map pre5              /* precedence 5 traffic */
  match access-group 105
class-map pre4              /* precedence 4 traffic */
  match access-group 183
class-map pre3              /* precedence 3 traffic */
  match access-group 184
class-map pre2              /* precedence 2 traffic */
  match access-group 185
class-map pre1              /* precedence 1 traffic */
  match access-group 186

policy-map cdf-pq                  /* Priority Queuing */
  class pre7                /* Robin PQ queue */
    priority 800
  class pre6
   bandwidth 100
   queue-limit 10
  class pre5
   bandwidth 100
   queue-limit 10
  class pre4
   bandwidth 100
   queue-limit 10
  class pre3
   bandwidth 100
   queue-limit 10
  class pre2
   bandwidth 100
   queue-limit 10
  class pre1
   bandwidth 100
   queue-limit 10
  class class-default
   bandwidth 100

interface FastEthernet0/0
 description test LAN
 ip address 192.168.72.1 255.255.255.0 secondary
 ip address 192.168.174.1 255.255.255.0 secondary
 ip address 192.168.184.1 255.255.255.0 secondary
 ip address 192.65.183.129 255.255.255.240 secondary
 ip address 192.168.186.1 255.255.255.0 secondary
 ip address 131.154.98.253 255.255.255.0 secondary
 ip address 192.168.190.14 255.255.255.248 secondary
 ip address 192.168.73.1 255.255.255.0
 no ip directed-broadcast
/* classification, marking and policing */
 rate-limit input access-group 150 1000000 8000 8000 conform-action set-prec-transmit 7
exceed-action drop   /* Robin traffic */
 rate-limit input access-group 161 1000000 8000 8000 conform-action set-prec-transmit 1
exceed-action set-prec-transmit 1
 rate-limit input access-group 162 1000000 8000 8000 conform-action set-prec-transmit 2
exceed-action set-prec-transmit 2
 rate-limit input access-group 163 1000000 8000 8000 conform-action set-prec-transmit 3
exceed-action set-prec-transmit 3
 rate-limit input access-group 164 1000000 8000 8000 conform-action set-prec-transmit 4
exceed-action set-prec-transmit 4
 rate-limit input access-group 165 1000000 8000 8000 conform-action set-prec-transmit 5
exceed-action set-prec-transmit 5
 rate-limit input access-group 166 1000000 8000 8000 conform-action set-prec-transmit 6
exceed-action set-prec-transmit 6
 ip route-cache policy
 load-interval 30
```

```
 full-duplex

interface ATM1/0.108 point-to-point
 description to ESnet (CDF QoS testing)
 ip address 192.168.190.2 255.255.255.252
 no ip directed-broadcast
 pvc 4/108
  tx-ring-limit 5
  service-policy output cdf-pq           /* scheduling is activated */
  vbr-nrt 2000 2000 1
  encapsulation aal5snap

access-list 150 permit ip any host 192.190.218.1
access-list 161 permit udp any host 192.190.218.2 eq 50001
access-list 162 permit udp any host 192.190.218.2 eq 50002
access-list 163 permit udp any host 192.190.218.2 eq 50003
access-list 164 permit udp any host 192.190.218.2 eq 50004
access-list 165 permit udp any host 192.190.218.2 eq 50005
access-list 166 permit tcp any host 192.190.218.2
access-list 180 permit ip any any precedence network
access-list 181 permit ip any any precedence internet
access-list 182 permit ip any any precedence critical
access-list 183 permit ip any any precedence flash-override
access-list 184 permit ip any any precedence flash
access-list 185 permit ip any any precedence immediate
access-list 186 permit ip any any precedence priority
```

APPENDIX D: The Production Network

ROUTING BETWEEN END NODES

```
 1  131.154.99.253 (131.154.99.253)  0.859 ms  0.800 ms  0.779 ms
 2  rc-cnaf.bo.garr.net (193.206.128.17)  2.161 ms  2.382 ms  3.582 ms
 3  rt-rc-2.bo.garr.net (193.206.134.157)  3.617 ms  2.199 ms  2.161 ms
 4  mi-bo-2.garr.net (193.206.134.5)  7.326 ms  8.867 ms  7.423 ms
 5  garr.ny3.ny.dante.net (212.1.200.17)  462.137 ms  123.842 ms 390.477 ms
 6  dante-gw.es.net (212.1.200.218)  115.315 ms  115.887 ms  123.528 ms
 7  chi-nyc.es.net (134.55.205.6)  136.922 ms  138.771 ms  136.147 ms
 8  fnal-chi.es.net (134.55.208.18)  162.010 ms  137.497 ms  139.997 ms
 9  f2.r-s-frw.fnal.gov (198.151.133.1)  140.192 ms  139.196 ms  142.722 ms
10  fe0-0-11.r-s-hub-fcc.fnal.gov (131.225.15.6)  141.933 ms  138.479 ms  151.367 ms
11  vme-qos-prod.fnal.gov (131.225.80.91)  141.931 ms  139.822 ms  138.087 ms
```

ROUND TRIP TIME (ping test)

From INFN CNAF (`robin1.cnaf.infn.it`) to FNAL (`vme-qos-prod.fnal.gov`)

```
--- vme-qos-prod.fnal.gov ping statistics ---
64 packets transmitted, 64 packets received, 0% packet loss
round-trip min/avg/max = 137.7/139.2/145.0 ms
```