SA WG2 Meeting #S2-129
8 - 12 October, 2018, Dongguan, China

S2-1810373

Using Priority Level for scheduling in 5G

Simulation study comparing Delay threshold Scheduling and Weighted Round Robin Scheduling in Mixed traffic scenarios.

Ericsson

© 3GPP 2018

Introduction

For LTE it is standardized that the Priority Level Parameter (PL) shall be used to prioritize Bearers with delay close to the Packet Delay Budget (PDB), and high PL.

For 5G, two main alternatives have been proposed, either that PL is defined similarly to LTE, or that PL is used to define the amount of resources that should be scheduled to each QoS Flow.

In this simulation study, two scheduling algorithms, based on the two PL definition alternatives are investigated:

- Prioritized Delay Threshold Scheduling – The LTE alternative
- Weighted Round Robin – Simple and popular scheduling algorithm.

© 3GPP 2018

Simulation details

- LTE system with 20 MHz bandwidth
  - Faster simulations than with 5G.
  - Results should be similar.
- 3GPP case1: 21 Cells, Cell radius=166 m
- Typical urban propagation

- Traffic is ftp download.
  - Large object downloads with TCP Cubic, on 5QI 8 (Best effort)
  - Small object download with TCP Cubic, on 5QI 8 and 5QI 5 (IMS)
  - Large object downloads with TCP Ledbat on 5QI 8
  - Large object downloads with TCP Cubic, on 5QI 6 (Prioritized)

© 3GPP 2018

Small&Large

Impact of scheduling algorithms on different sized data objects is investigated in a mixed traffic scenario.

- Traffic consists of small and large ftp downloads.
  - Small:100 kbyte, Large: 10 Mbyte
- TCP Cubic used for both traffic types.
- Load from large downloads is increased until subband utilization is 90%.
- Both traffic types are using 5QI 8 (PDB=300ms)
- Two scheduling algorithms:
  - Round Robin (RR)
  - Delay Threshold Sched. (Del)
    - Data over delay threshold get absolute priority.
    - Delay Threshold: PDB-20 ms
- Figures at right show channel utilization and cell throughput. The cell throughput is lower with delay threshold scheduling, since users with high delay and bad radio conditions are prioritized.

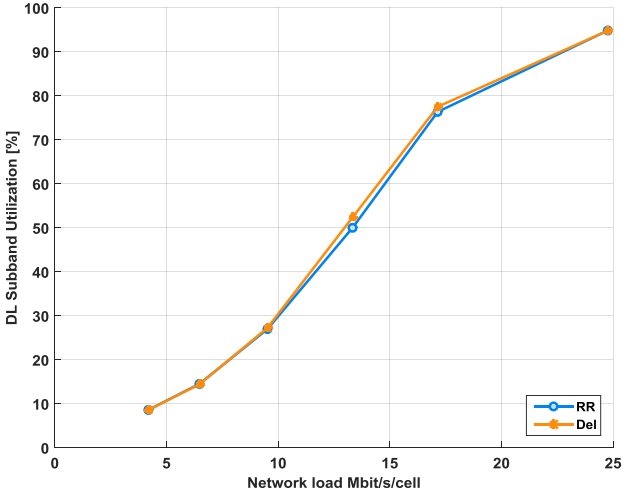

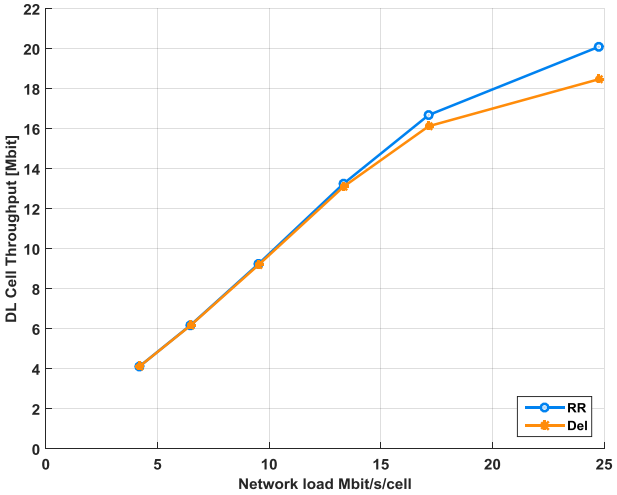

© 3GPP 2018

Small&Large -Bitrate

Average and 5th percentile Object bitrate is shown for small and large downloads, using both scheduling algorithms.

- Bitrates of large objects are much higher, since TCP slowstart is dominating small object transmission.
- Bitrates of Small objects are lower with delay threshold scheduling. Especially the 5% with worst quality are penalized by the algorithm.
- Average bitrates of Large objects are slightly higher with delay threshold scheduling.

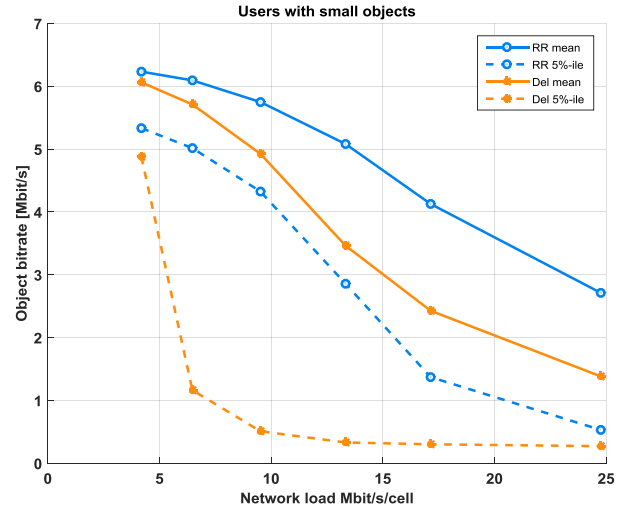

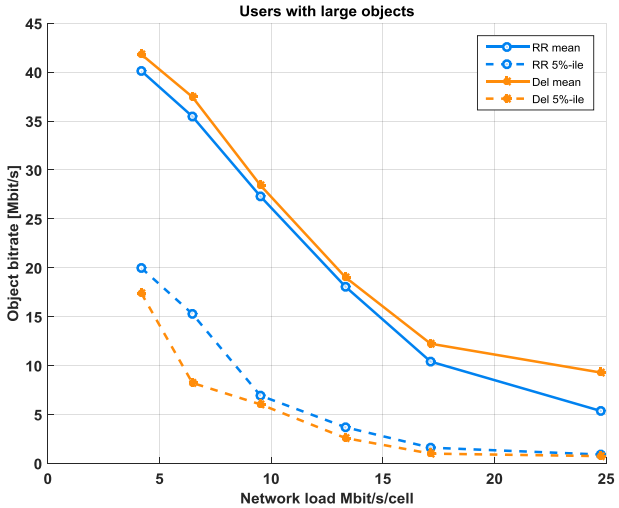

Small objects will get lower bitrates due to TCP effects. Delay Threshold Scheduling magnifies the effect. Round Robin is more fair.

© 3GPP 2018

Small&Large – Delay jitter

- The QoS Packet Delay target is that the 98th percentile packets shall have delays<PDB. Therefore, the delay jitter of the objects was defined as the 98th percentile segment delay.
- The figures shows the average and 95th percentile delay jitter for small and large objects.
- Again, performance of small objects are much better with RR. (With exception of the highest load case.)
- Delay jitter of the large objects is lower with Delay threshold scheduling, but even for average objects, it is above the QoS target of 300 ms.
- When the packet delay is above 300 ms, AQM drops will occur, and TCP Cubic will limit the rate.

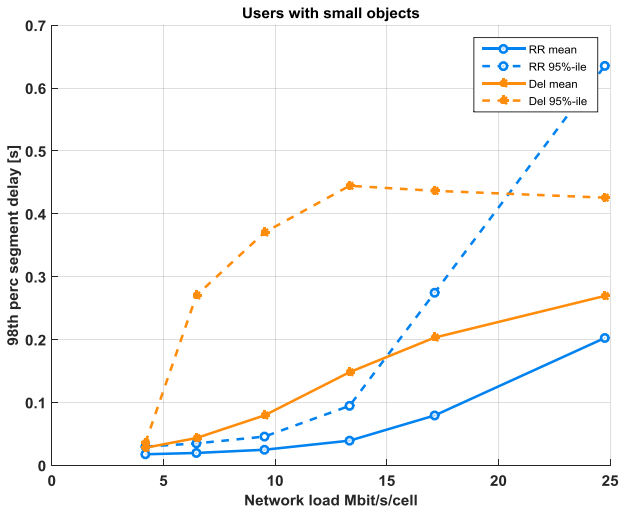

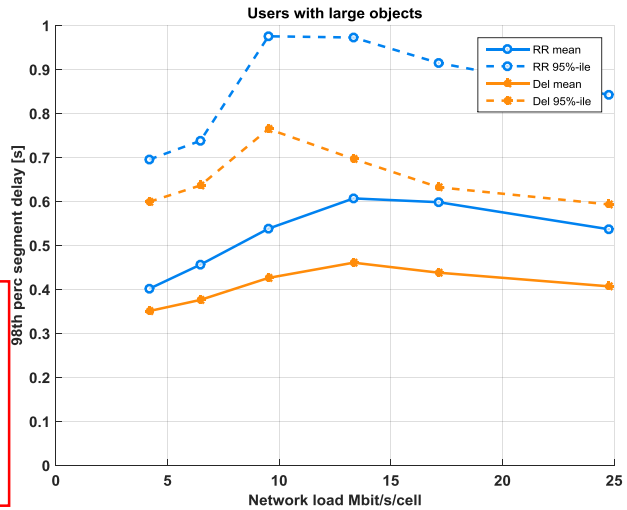

Small objects get higher delay jitter with delay threshold scheduling.
For large objects, delay jitter is lower with delay threshold scheduling, but it is above the target 300 ms.

© 3GPP 2018

Large Objects - Delay distribution

Figure shows the delay distribution for transmission of large objects in three load levels. (10%, 30% and 80% channel utilization)

- For both scheduling methods and all loads, there is no major difference.
  - About 70-85% of packets delays are <300 ms.
  - The median delay is approx. 230 ms

At delays of 300 ms, AQM will start dropping some packets.

- The packets are so large, so that even in low load, the transmission will use all available capacity, so TCP congestion control is needed to avoid congestion. TCP Cubic is loss-based and will only react on lost packets.
- Cubic will adapt bitrate so that median delays are below AQM threshold.
- Some delays will be significantly higher than AQM threshold, due to changed load conditions. (This is especially true in medium load, where the cell load will change between 1 and 2 users frequently.)

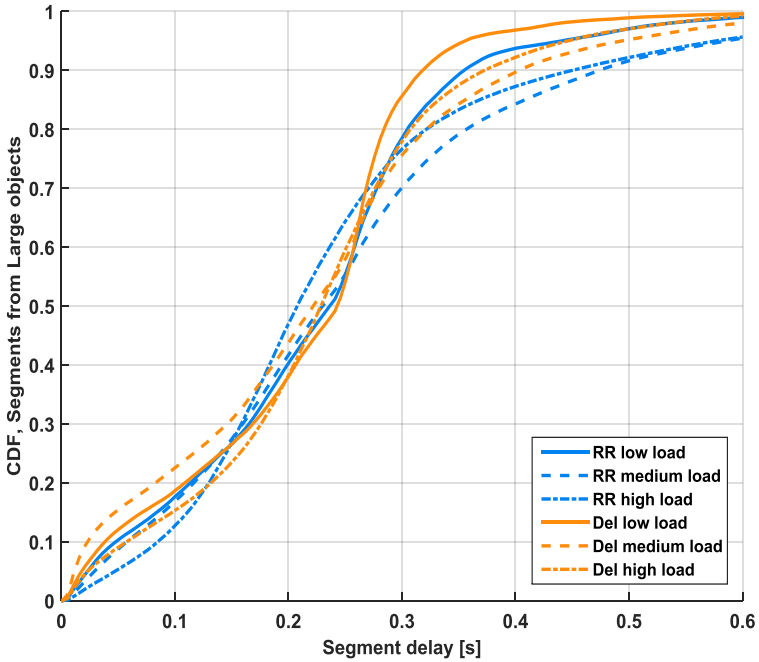

Scheduling have small impact on the delay for large transmissions with TCP Cubic.

© 3GPP 2018

Small IMS Objects

In order to investigate performance of IMS traffic, the traffic scenario was modified, so that the small objects are part of an IMS Flow, sent on 5QI 5. The large objects are still sent on 5QI 6.

- 5QI 5: PDB=100 ms, PL=10. 5QI 6: PDB=300 ms, PL=60.
- Three scheduling algorithms:
  - Round Robin (RR)
    - IMS not prioritized
  - Weighted RR (WRR)
    - The scheduling weight is W*RR weight.
    - IMS: W=5 Large objects: W=1
  - Prioritized Delay Threshold Scheduling (DelPrio)
    - IMS traffic over delay threshold have absolute priority over other traffic
    - Delay Threshold = PDB-20 ms

© 3GPP 2018

Small IMS Objects

- The object bitrate is highest with WRR scheduling.
- Prioritized delay scheduling perform even worse than unprioritized RR at low and medium load.
- The delay jitter (98th percentile delay) is lowest with WRR scheduling.
- With RR scheduling, the delay jitter is low at low load, but increase quickly when the load increase.
- With DelPrio scheduling, the delay jitter of the 95th percentile users is just above the PDB for most load cases. With a larger delay margin, the algorithm might work, but the delays are longer than with WRR.

For IMS data, WRR scheduling outperforms delPrio Scheduling in this traffic scenario.

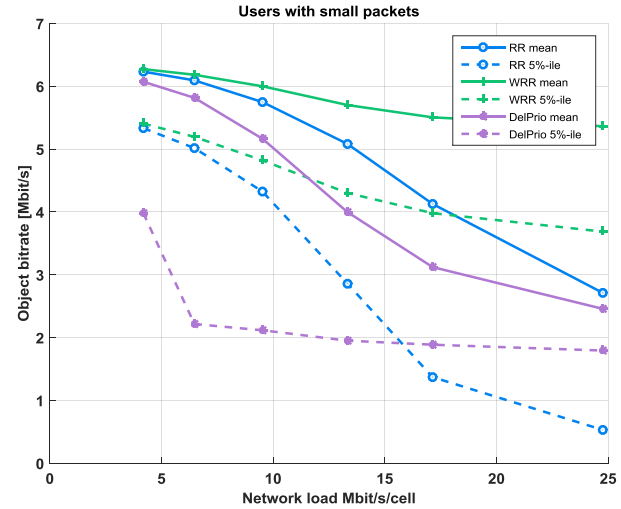

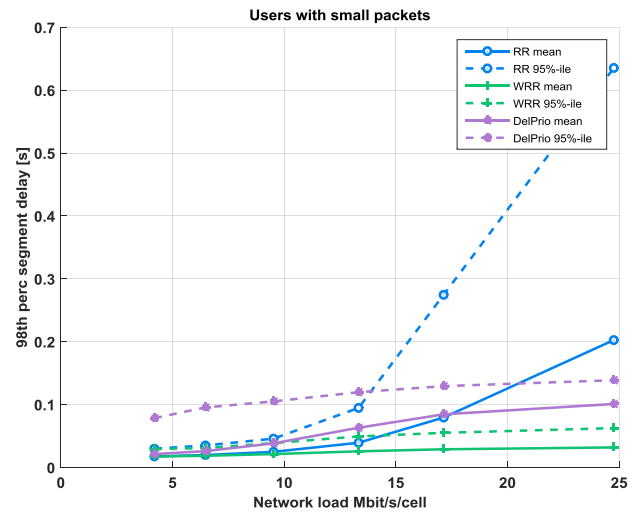

© 3GPP 2018

Small IMS Objects – Delay distribution

Figure shows the delay distribution for transmission of small prioritized objects in three load levels, and RR vs DelPrio scheduling.

- With Round Robin without priority:
  - Short delays at low and medium load
  - Some very long delays at high load
  - Small data flows are given the same amount of resources as large flows, so the throughput will be sufficient for small IMS objects, even without priority, unless the cell is very heavily loaded.
- With Prioritized Delay scheduling:
  - Short delays when the cell is empty.
  - Delay > delay threshold (80ms) when there is other traffic present.
  - A fraction of packets have delays over 100 ms, even at medium load.
  - No packets with delays over 200 ms.

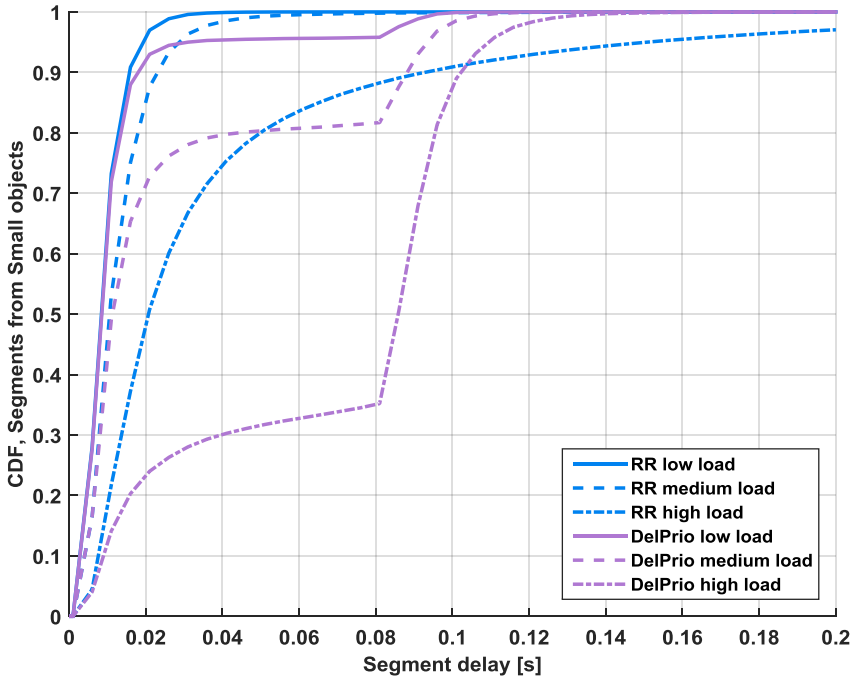

RR scheduling outperforms delPrio even without prioritization for most load cases.

© 3GPP 2018

TCP Ledbat

TCP Ledbat is a TCP congestion control algorithm which is using delay measurement to avoid network congestion and limit transmission delay. It is popular for large file transfers.

- Figures shows results when the whole network is filled with traffic using the same TCP congestion algorithm.
- Objects with TCP Ledbat have slightly lower bitrates than TCP Cubic, but the 98 percentile delay is much lower. (The 95th percentile Ledbat objects have lower jitter than the average Cubic object.)
- With both algorithms, congestion control works well. The delay jitter does not increase at high load.

TCP Ledbat can be used by applications to limit transmission delays.

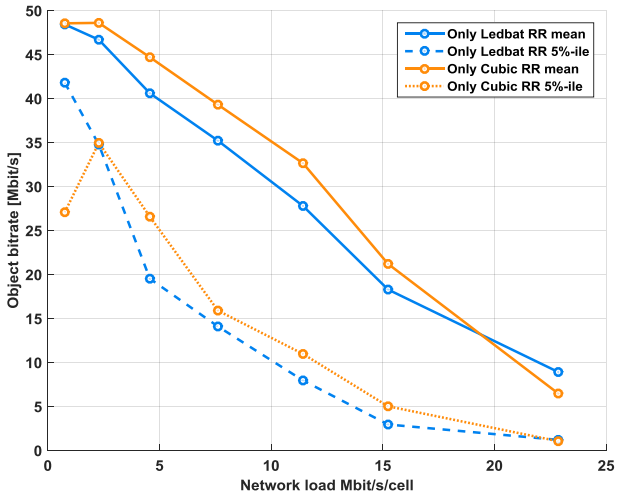

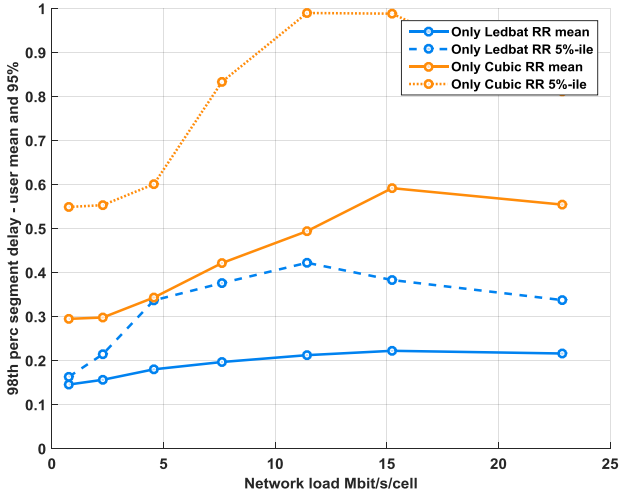

© 3GPP 2018

TCP Ledbat – Mixed traffic

We investigated a mixed traffic scenario where the network is shared by TCP Ledbat traffic and TCP Cubic.

- TCP Cubic traffic is increased until channel utilization is over 90%.
- Both RR and Delay Threshold Scheduling is tested.
- The bitrates are quite similar with RR (green), but with Del scheduling (purple), the Ledbat traffic (triangles) have lower rates.
- Using RR scheduling (green), the delay jitter is much lower with Ledbat (triangles) than Cubic (circles).
- With Del scheduling (purple), the delay jitter of Ledbat (triangles) are still lower, but it is increasing when load increase. – Cubic objects are prioritized, since they have higher delays.

Delay Threshold Scheduling will prioritize flows with large delays, and sabotage TCP Ledbat's attempts to limit delay.
Round Robin will distribute resources fairly.

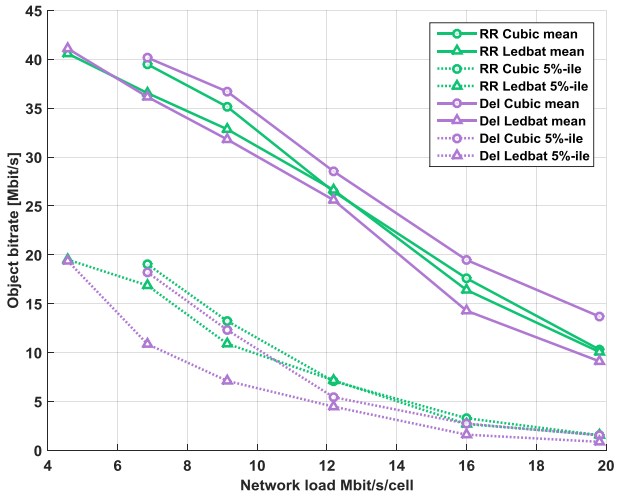

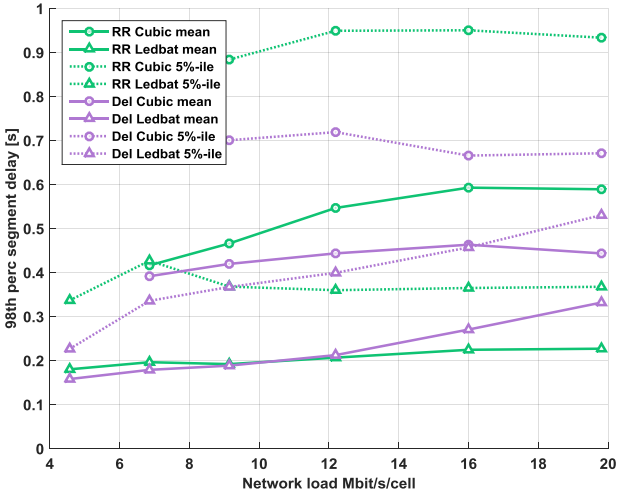

© 3GPP 2018

High&Low Priority

The impact of prioritization is investigated.

- Low prio objects sent on 5QI 8
- High prio objects sent on 5QI 6
- Same PDB, but different Priority Level.
- All Objects are large, using TCP Cubic.

- With WRR (red) the high prio objects (triangles) have higher mean and 5th percentile bitrates.
- With DelPrio (yellow) there is very little differentiation, except at the highest load.

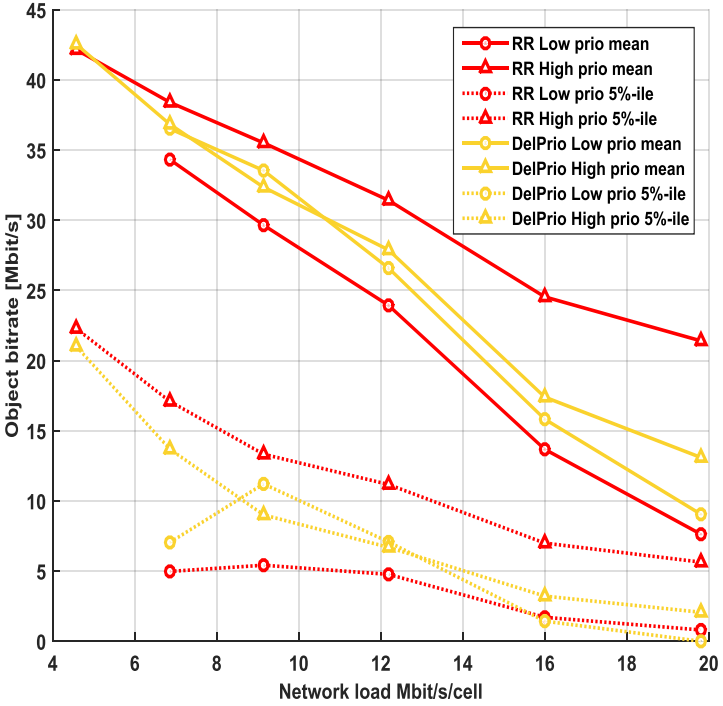

Weighted Round Robin will distribute more resources to prioritized flows, as soon as there are scheduling delays.
Delay Threshold Scheduling will only give more resources to prioritized flows when there are several flows with delays close to PDB. – Less differentiation

© 3GPP 2018

Conclusions

Scheduling cannot be used to limit delays for large transmissions with TCP Cubic. Cubic will adapt the data rate to the AQM threshold.

Scheduling will affect short transmissions, and prioritized transmissions. In all investigated scenarios, Weighted Round Robin performed better than Delay Threshold Scheduling.

With Delay Threshold Scheduling:

- The worst-case delays are limited at extreme congestion for flows that are prioritized, and not greedy.
- Large and Greedy flows will be prioritized over:
  - Flow with little data
  - Prioritized flows with delays < PDB
  - Flows that adapt the rate to limit delay. with conservative TCP CC like TCP Ledbat.

With Weighted Round Robin we will get:

- Fair resource distribution between flows.
- Low delays for flows with little data.
- Predictable behavior for CC algorithms based on transmission delay.

For large TCP flows, the scheduling should be Resource based in order to provide a fair distribution of radio resources. However, delay threshold scheduling does limit delay jitter at high loads, so hybrid solutions might also work well, where some resources are scheduled based on delay thresholds.

© 3GPP 2018