IETF89 AQM Evaluation guidelines
draft-kuhn-aqm-eval-guidelines-00

Nicolas Kuhn [1] Preethi Natarajan [2] David Ros [3] Naeem Khademi [4]

[1]IMT Telecom Bretagne, France

[2]Cisco Systems, United States

[3]Simula Research Laboratory, Norway

[4]University of Oslo, Norway

Contributors

- S. Akhtar
- A.B. Bagayoko
- F. Baker
- D. Collier-Brown
- G. Fairhurst
- T. Hoiland-Jorgensen
- C. Kulatunga
- R. Pan
- D. Taht
- M. Welzl

Table of contents

Context

- Active Queue Management (AQM) addresses the concerns arising from using unnecessarily large and unmanaged buffers
- AQM : one of the solutions to the bufferbloat issues (and reduce end-to-end latency)
- how do we get the confidence that a specific AQM is better than drop tail and thus "safe" to deploy ?
- the WG requires guidelines to ascertain whether the WG should undertake an AQM proposal

Objectives

- define operating regions of an AQM proposal and discuss the parameters sensitivity
- evaluation guidelines for very diverse network environments
- guidelines for performance evaluation :
  - how does the proposal compare drop-tail :
  - trade off between reducing the latency and maximizing the goodput
- guidelines for safe deployment :
  - how safe is it to deploy the proposal compared to the issues encountered by RED

Table of contents

Metrics to evaluate the trade-off between latency and goodput

- provide generic metrics that can be exploited to evaluate the performance of an AQM, whatever the context
- propose set of metrics to effectively evaluate the trade-off between latency and goodput
- an AQM proposal should :
  - control latency at a desired level
  - minimize the hit on goodput

Queue-related metrics

- Why queue-related metrics ?
  - queue-related metrics MUST be considered to understand the behavior of an AQM and the possible impact of its internal parameters
- Which queue-related metrics ?
  - link utilization : [RFC5136] "the utilization now represents the fraction of the capacity that is being used and is a value between zero (meaning nothing is used) and one (meaning the link is fully saturated)"
  - queuing delay and queue size : The queuing delay is the time a packet waits in a queue until it can be forwarded to the lower layers. The queue size is the number of bytes which are occupying the queue.
  - two classes of packet loss : AQM-induced losses and buffer overflow
    - long-term packet loss probability
    - time interval between consecutive losses
  - any other queue-related metrics for performance evaluation ?

End-to-end metrics

- Why end-to-end metrics ?
  - End-to-end metrics MUST be considered to evaluate the consequences of introducing an AQM on the latency and the goodput
- Which end-to-end metrics ?
  - flow completion time : distribution of the flow completion time depending on the flow size
  - packet loss : long term packet loss probability, loss inter-arrival time and packet loss pattern
  - packet loss synchronization : degree of synchronization of loss events between two flows on the same path
  - goodput : important if scheduling comes into play (possible flow starvation) - end-to-end assessment of how well the AQM improves transport application performance
  - latency and jitter : differs from queuing delay and depends on traffic and topology
  - QoE : metrics to assess AQM's performance for dedicated context (VoIP, video-streaming, data centers, web)
  - any other end-to-end metrics for performance evaluation ?

Table of contents

Evaluation scenarios

- provide set of scenarios that could be considered to evaluate the AQM performance
- each scenario is a potential operating region for the tested algorithm
- an evaluation report on a given AQM should make it clear when the parameters of an AQM had to be externally adjusted for the AQM to perform on various scenarios
- for each scenario, the tester should exploit the metrics presented earlier
- the evaluation scenarios can be divided into two classes :
  - generic scenarios that must be considered
  - specific network environments scenarios that could be considered to evaluate the performance of an AQM in particular conditions
- objective of this presentation :
  - justify and detail the scenarios
  - discuss which scenario MUST and which MAY be considered
  - agree on a structure for the performance evaluation
  - detailed topology or traffic modelisation may not be discussed here, but comments are welcome on the aqm-list

Methodology

- sufficiently detailed description of the test setup should be provided (that would allow other to replicate the tests)
- the test setup MAY include software and hardware versions (in case they impact on the AQM performance).
- the tester MAY make its data available ?
- the guidelines are not bound to a particular evaluation tool set, however :
  - proposals SHOULD (MUST ?) be experimented on real systems ;
  - proposals MAY be evaluated with event-driven simulations (such as NS-2, NS-3, OMNET, etc.).

Topology

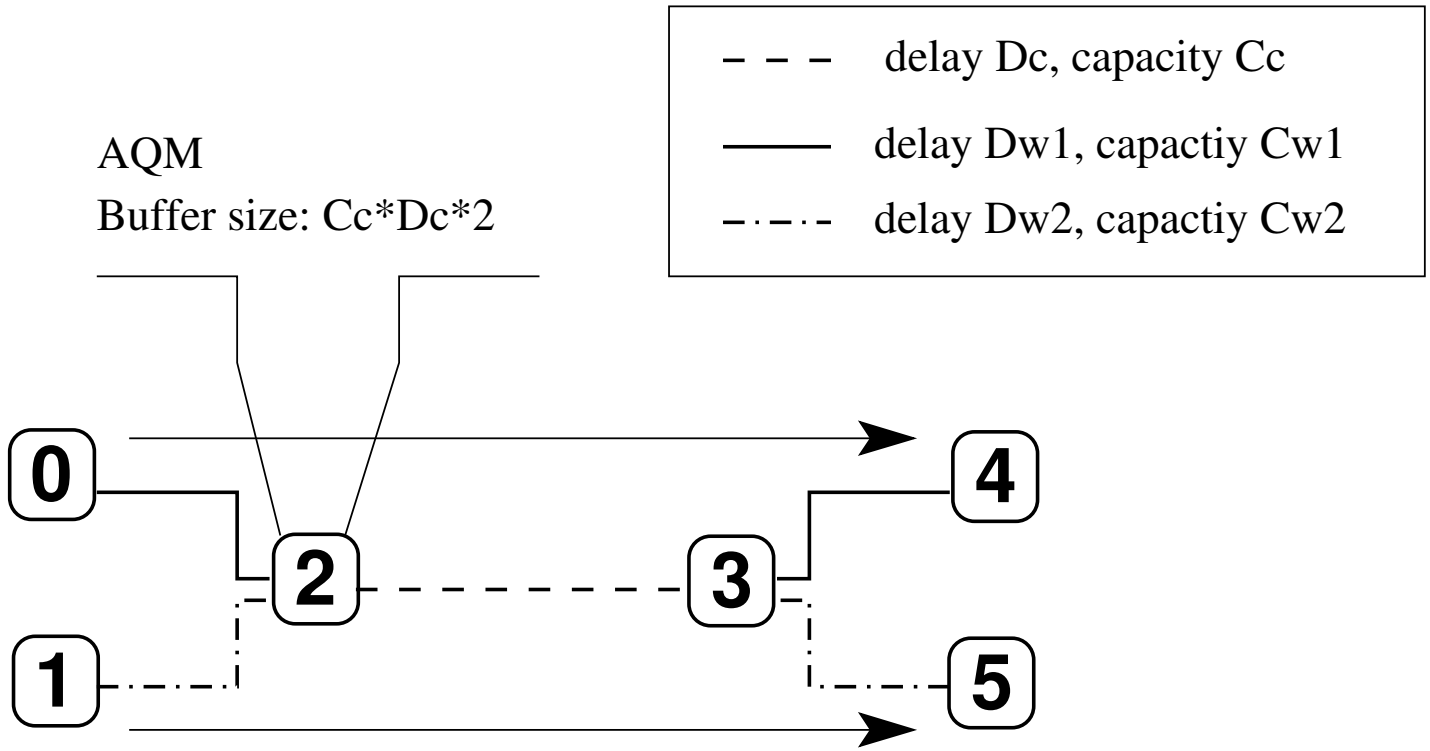

- classical dumbbell network
- AQM in the router before the bottleneck
- the (RTT,Capacity) of each link is independent from the others
- size of the buffer should be carefully set, considering the bandwidth-delay product

Generic scenarios - Traffic profiles

- Why this scenario ?
  - applications run over different flavors of TCP (unresponsive flows (UDP ...), or aggressive flows)
  - AQM should ensure queuing delay is under control with traffic profiles
- Traffic profiles

Case 1	TCP-friendly Sender	TCP New Reno or others
Case 2	aggressive Transport Sender	TCP Cubic or others
Case 3	unresponsive Transport Sender	UDP flows
Case 4	TCP inigial congestion window	mix of TCP New Reno, TCP Cubic with IW3 and IW10
Case 5	traffic mix	TCP transfer, HTTP traffic, VoIP Gaming, CBR, adaptive video streaming

- Output
  - (at the AQM level) queuing delay vs. link utilization vs. drop rate
  - (e2e level) end-to-end delay vs. goodput vs. drop rate

Generic scenarios - Burst absorption

- Why this scenario ?
  - Bursty packet arrivals
  - queuing delay spikes must be minimized
  - performance penalties (losses) for ongoing flows must be minimized
- Traffic profiles
  - Bursty traffic :

Generic	one CBR (UDP) traffic one repeating TCP transfer burst of packets (with various sizes of burst)
Realistic	one CBR (UDP) traffic repeating TCP transfer (IW10) HTTP web traffic bursty video frames

- Output
  - (at the AQM level) queuing delay vs. link utilization vs. drop rate
  - (e2e level) end-to-end delay vs. goodput vs. drop rate
  - flow completion time
  - Jitter, latency

Generic scenarios - Inter-RTT and intra-protocol fairness

- Why this scenario ?
  - asymmetry (various RTT) SHOULD be considered : fairness between the flows
  - AQM must be evaluated against a set of RTT
- Topology
  - To evaluate inter-RTT fairness :
    - flows with RTT in [5ms ;100ms]
    - flows with RTT in [100ms ;200ms]
  - To evaluate intra-protocol fairness :
    - flows with RTT in [5ms ;200ms]
- Output
  - (at the AQM level) queuing delay vs. link utilization vs. drop rate
  - (e2e level) end-to-end delay vs. goodput vs. drop rate

Generic scenarios - Fluctuating network conditions

- Why this scenario ?
  - varying operating conditions (time of day or deployment scenario)
  - stability of the AQM's parameters over time is challenged
- Traffic profiles
  - mild congestion
  - medium congestion
  - heavy congestion
  - varying available bandwidth
- Output
  - (at the AQM level) queuing delay vs. link utilization vs. drop rate
  - (e2e level) end-to-end delay vs. goodput vs. drop rate

Diverse network environments - Medium bandwidth, medium delay

- Why this scenario ?
  - this scenario is introduced to carefully evaluate AQM proposals in a generic context
  - the tester COULD consider this use case to define the operating region of the AQM
- Toopology and Traffic profiles

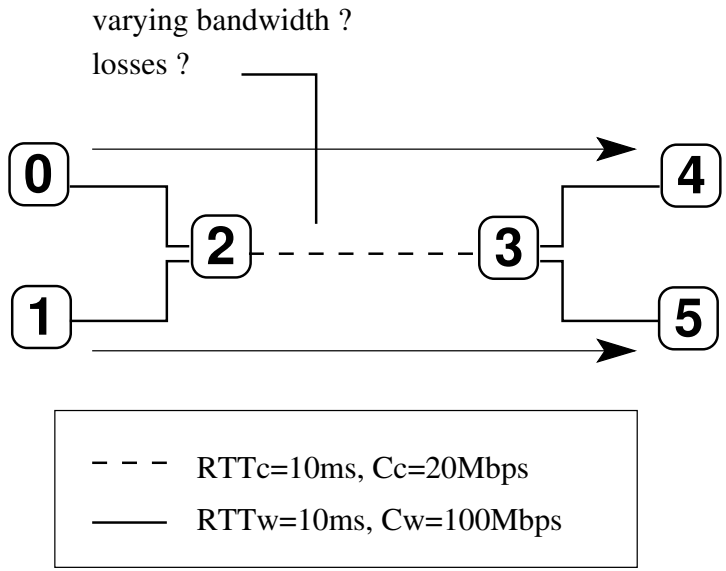

  - Traffic :
    - repeating TCP transfers
    - continuous TCP transfer
    - HTTP web traffic
- Output
  - (at the AQM level) queuing delay vs. link utilization vs. drop rate
  - (e2e level) end-to-end delay vs. goodput vs. drop rate

Diverse network environments - Low bandwidth, high delay

- Why this scenario ?
  - low bandwidth and high delay : seriously challenged burst absorption capacity
  - this use case : operating region of the AQM
- Toopology and Traffic profiles

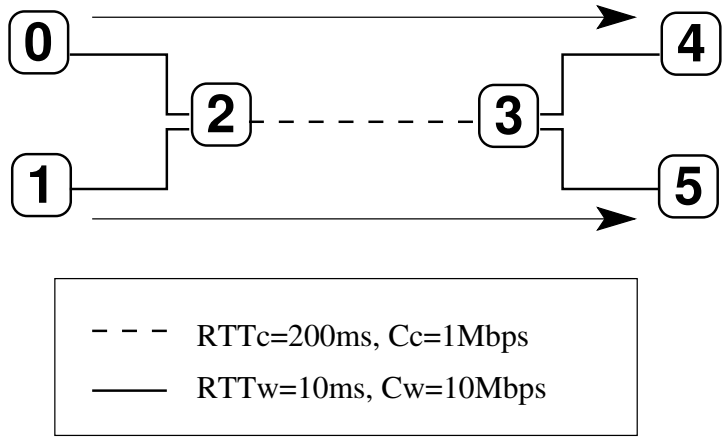

  - Traffic :
    - repeating TCP transfers
    - continuous TCP transfer
    - HTTP web traffic
- Output
  - (at the AQM level) queuing delay vs. link utilization vs. drop rate
  - (e2e level) end-to-end delay vs. goodput vs. drop rate

Diverse network environments - High bandwidth, low delay

- Why this scenario ?
  - high bandwidth and low delay : may require updated thresholds, auto-tuning of an AQM is challenged
  - this use case : define an operating region of the AQM
- Toopology and Traffic profiles

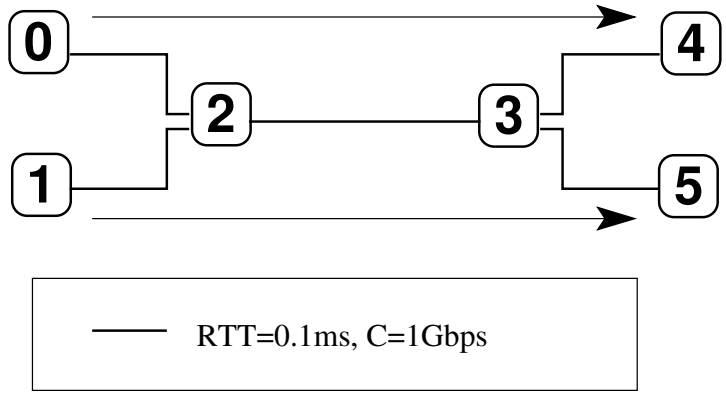

  - Traffic :
    - repeating TCP transfers
- Output
  - (at the AQM level) queuing delay vs. link utilization vs. drop rate
  - (e2e level) end-to-end delay vs. goodput vs. drop rate

Diverse network environments - Small and large buffers

- Why this scenario ?
  - size of the buffers impacts on AQMs performance (even if based on queue length or queueing delay)
  - low buffer (i.e. 1/10 BDP) and large buffer (i.e. 10 BDP)
- Toopology and Traffic profiles

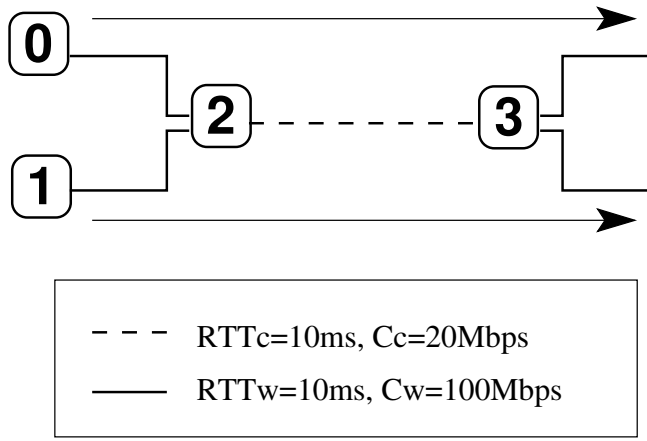

  - Traffic :
    - repeating TCP transfers
    - continuous TCP transfer
    - HTTP web traffic
- Output
  - (at the AQM level) queuing delay vs. link utilization vs. drop rate
  - (e2e level) end-to-end delay vs. goodput vs. drop rate

Performance evaluation : are we missing anything ?

- Scenarios that MUST be considered :
  - various traffic profiles (unresponsive flows, aggressive flows, etc.)
  - burst absorption capacity
  - inter-RTT and intra-protocol fairness
  - fluctuating network conditions
- Scenarios that MAY be considered :
  - medium bandwidth, medium delay
  - low bandwidth, high delay
  - high bandwidth, low delay
  - small and large buffers
- are we missing anything ?

Table of contents

Deployment

- This section details deployment issues that MUST be discussed (such as stability, implementation cost, implementation feasibility, control knob)
- This section helps to discuss how safe is it deploy the proposal compared to the issues encountered by RED

Operator control knobs and auto-tuning

- Requirements
  - an AQM scheme should be stable in varying conditions without the need for external tuning (employing auto-tuning if needed)
  - an AQM scheme should minimize the control knobs exposed for operator tuning, to be more user-friendly and easier to deploy and debug

Parameter sensitivity and stability analysis

- Requirements
  - AQM proposals SHOULD (MAY ?) provide background material to discuss its stability
  - AQM proposals SHOULD (MAY ?) provide the input parameter space within the AQM operates as expected
  - for parameters that are auto-tuned, the material SHOULD (MAY ?) include stability analysis of the auto-tuning mechanism(s) as well
  - the impact of every externally tuned parameter MUST be discussed
  - these guidelines discourage unnecesseray external tuning

Implementation cost

- Requirements that help identify costs associated with implementing the AQM on a particular hardware or software platform
  - AQM proposals SHOULD provide pseudo-code for the complete AQM scheme, highlighting generic implementation
  - AQM proposals SHOULD highlight parts of AQM logic that are platform dependent and discuss if and how AQM behavior could be impacted by the platform

Interaction with packet scheduling

- Requirements
  - the tester MUST discuss the feasibility to add scheduling on top of its algorithm
  - this discussion MAY detail if dropping policy is placed while packets are enqued and dequed

ECN behavior

- Requirements
  - An AQM scheme SHOULD support ECN
  - An AQM SHOULD leverage ECN as an initial means to control queuing delay before resorting to packet drops
  - An AQM scheme SHOULD self-adapt and remain stable even with faulty and/or unresponsive ECN implementations

Packet sizes and congestion notification

- Requirements
  - An AQM scheme SHOULD adhere to recommendations outlined in [I-D.ietf-tsvwg-byte-pkt-congest]
  - SHOULD NOT provide undue advantage to flows with smaller packets

Deployment : are we missing anything ?

- operator control knobs and auto-tuning
- parameter sensitivity and stability analysis
- implementation cost
- interaction with packet scheduling
- ECN behavior
- packet sizes and congestion notification
- are we missing anything ?

Table of contents

Comparing AQMs

- the guidelines mentioned above may be used for comparing AQMs
- this memo recommends that AQM schemes MUST be compared against both performance and deployment categories

Comparing AQMs

- Performance
  - AQM schemes MUST be compared against all the generic scenarios
  - AQM schemes MAY be compared for specific network environments
  - if an AQM scheme's parameter(s) were externally tuned for optimization or other purposes, these values MUST be disclosed
  - Fair comparison of AQM schemes :
    - Problem 1 : AQM schemes belong to different varieties such as queue-length based scheme (ex :RED) or queue-delay based scheme (ex : CoDel, PIE)
    - Recommendation : Identify comparable control parameters and comparable input values (ex : qlen and target_delay)
    - Problem 2 : AQM schemes expose different control knobs associated with dierent semantics (ex :CoDel's "queueing delay target" and PIE's "queueing delay reference" are different
    - Recommendation : Compare over a range of input parameters (ex : 5ms, 10ms, 15ms target delay values)
- Deployment
  - all the deployment criteria discussed earlier must be considered