Measurement-Based Admission Control for Flow-Aware Implicit Service Differentiation

Jordan Augé (UPMC Sorbonnes Universités & CNRS),
Sara Oueslati (Orange Labs),
James Roberts (INRIA)

23rd International Teletraffic Congress (ITC 2011), San Francisco (CA), Sep 6-8, 2011

Outline

Outline

Introduction: the Cross-Protect (XP) router

A proposition to realize implicit differentiation of streaming and elastic flows, and guarantee their performance.

Introduction: the Cross-Protect (XP) router

A proposition to realize implicit differentiation of streaming and elastic flows, and guarantee their performance.

Streaming and elastic flows

Introduction: the Cross-Protect (XP) router

A proposition to realize implicit differentiation of streaming and elastic flows, and guarantee their performance.

Streaming and elastic flows

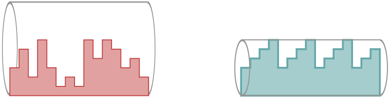

Introduction: the Cross-Protect (XP) router

A proposition to realize implicit differentiation of streaming and elastic flows, and guarantee their performance.

Streaming and elastic flows

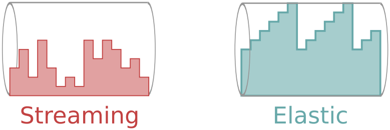

Introduction: the Cross-Protect (XP) router

A proposition to realize implicit differentiation of streaming and elastic flows, and guarantee their performance.

Streaming and elastic flows

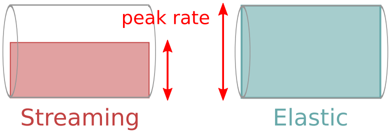

Introduction: the Cross-Protect (XP) router

A proposition to realize implicit differentiation of streaming and elastic flows, and guarantee their performance.

Streaming and elastic flows

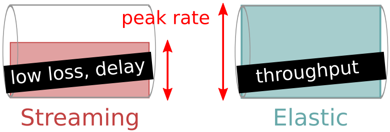

Introduction: the Cross-Protect (XP) router

A proposition to realize implicit differentiation of streaming and elastic flows, and guarantee their performance.

Streaming and elastic flows

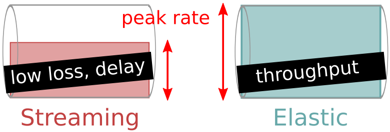

Combination of two mechanisms

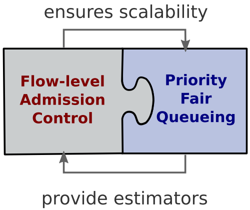

Introduction: the Cross-Protect (XP) router

A proposition to realize implicit differentiation of streaming and elastic flows, and guarantee their performance.

Streaming and elastic flows

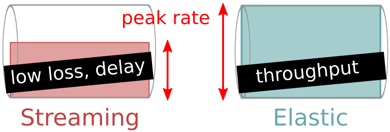

Combination of two mechanisms

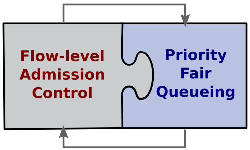

Assumption for streaming traffic

Peak rate r less than a threshold p (e.g. $p = .01C$)

- licit in backbone networks

example: 16Mb/s HDTV flows on a 40Gb/s link ($\sim .001C$)

Motivations: link operating regimes

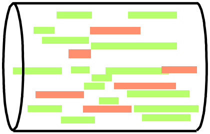

Motivations: link operating regimes

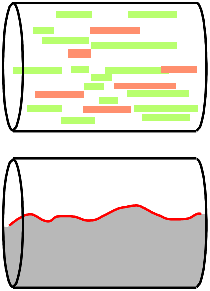

Motivations: link operating regimes

"transparent"

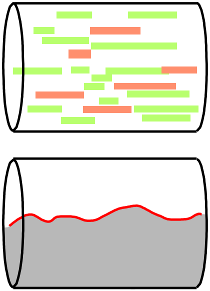

negligible loss
and delay

FIFO sufficient

Motivations: link operating regimes

"transparent" | "elastic"

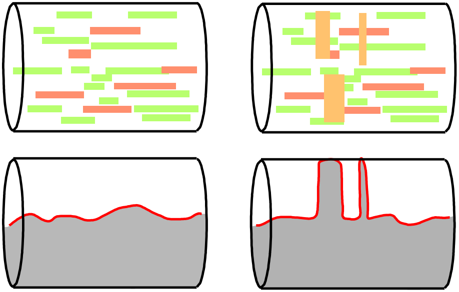

negligible loss and delay

FIFO sufficient

excellent for elastic, some streaming loss

needs differenciation

Motivations: link operating regimes

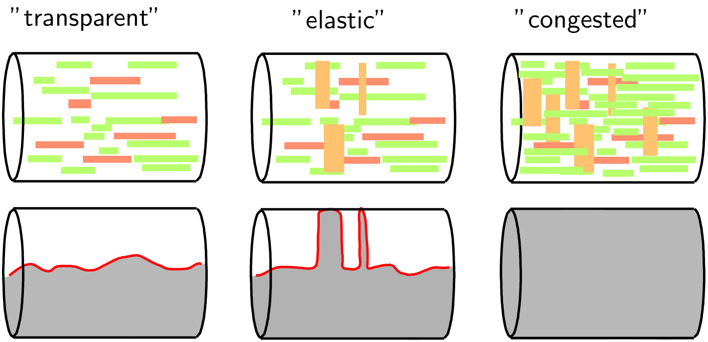

negligible loss and delay

FIFO sufficient

excellent for elastic, some streaming loss

needs differenciation

low throughput, significant loss

needs overload control

Priority Fair Queueing

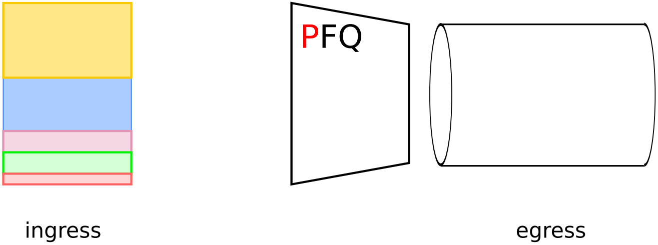

Priority Fair Queueing

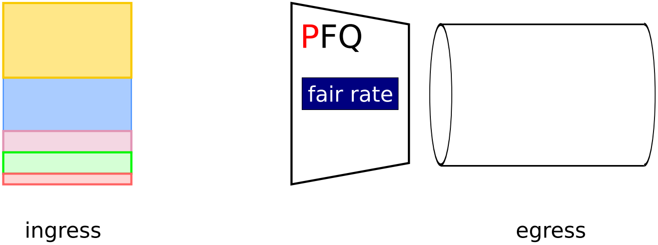

Priority Fair Queueing

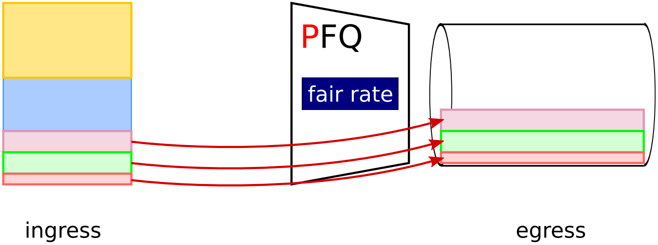

Priority Fair Queueing

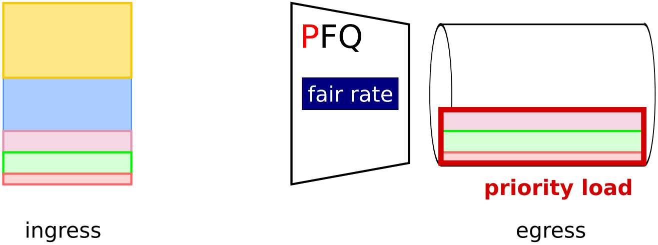

Priority Fair Queueing

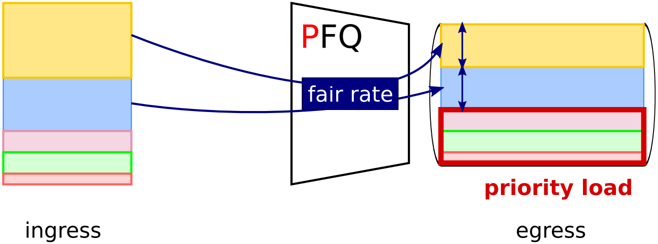

Priority Fair Queueing

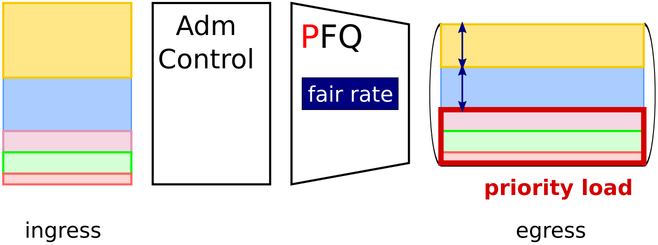

Priority Fair Queueing

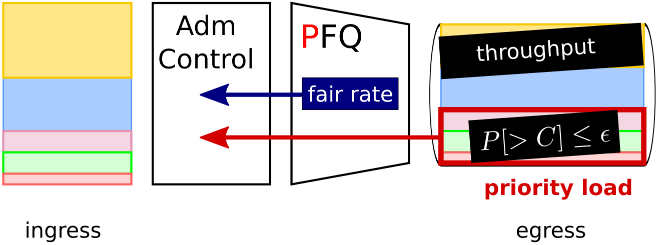

Priority Fair Queueing

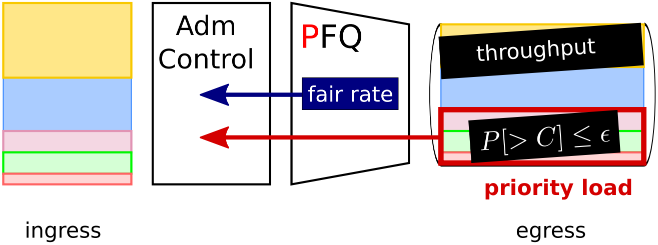

Discrimination: flows under/over fair rate

Bufferless multiplexing context for streaming

- Aggregate load
- Flow peak rate

Performance insensitive to detailed traffic characteristics

Outline

Requirements for a Cross-Protect MBAC

Discriminate flows

- Assumption: streaming flows peak rate $< p$
- Flows over p typically elastic and/or should be adaptive.
- Both will be handled in the priority queue

Cope with a minimal set of assumptions

- priority load and fair rate estimates
- maximum protected flow peak rate p
- NO indication of end of flows (timeout)

Timescale Decomposition Approach [Grossglauser,Tse]

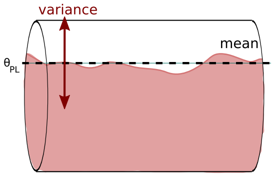

Approach:

Spare bandwidth to prevent overload:

Timescale Decomposition Approach [Grossglauser,Tse]

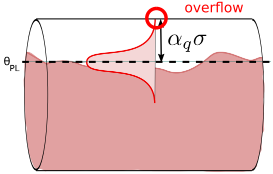

Approach:

Spare bandwidth to prevent overload:

- Gaussian approximation of aggregate load

Timescale Decomposition Approach [Grossglauser,Tse]

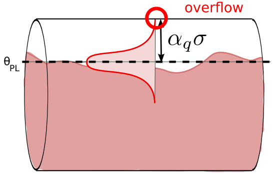

Approach:

Spare bandwidth to prevent overload:

- Gaussian approximation of aggregate load
- Introduction of a critical timescale
  $\sim$ flow timescale

Timescale Decomposition Approach [Grossglauser,Tse]

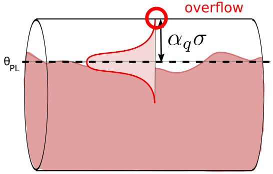

Approach:

Spare bandwidth to prevent overload:

- Gaussian approximation of aggregate load
- Introduction of a critical timescale
  $\sim$ flow timescale

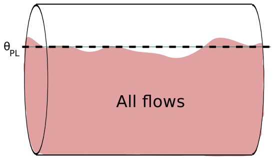

Limitations

- Threshold set for all traffic
- Prevent favourable state for differentiation

Timescale Decomposition Approach [Grossglauser,Tse]

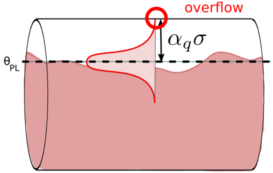

Approach:

Spare bandwidth to prevent overload:

- Gaussian approximation of aggregate load
- Introduction of a critical timescale $\sim$ flow timescale

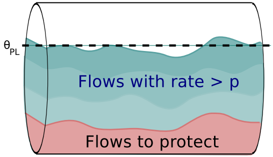

Limitations

- Threshold set for all traffic
- Prevent favourable state for differentiation

Timescale Decomposition Approach [Grossglauser,Tse]

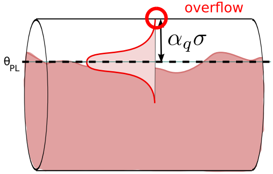

Approach:

Spare bandwidth to prevent overload:

- Gaussian approximation of aggregate load
- Introduction of a critical timescale $\sim$ flow timescale

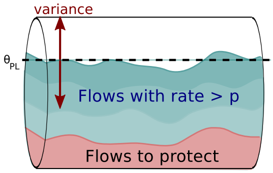

Limitations

- Threshold set for all traffic
- Prevent favourable state for differentiation

A Poisson approximation (1/2)

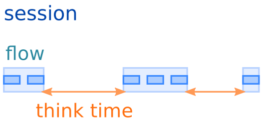

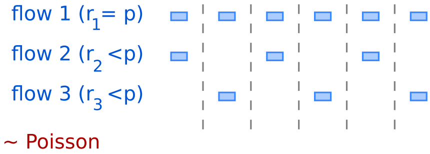

A Poisson approximation (1/2)

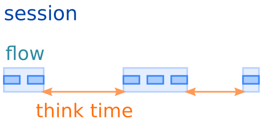

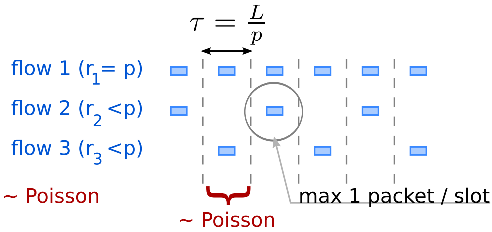

A Poisson approximation (1/2)

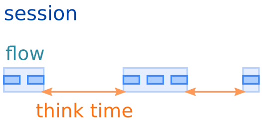

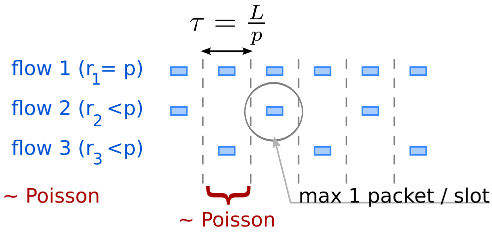

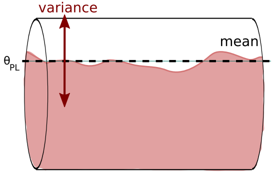

A_t: smoothed priority load.

$$\widehat{\sigma}_t^2 = A_t p$$ **(Poisson)**

79% utilization for $p = .01C$ and overflow prob. $\epsilon = 10^{-2}$

A Poisson approximation (2/2)

Flows of peak rate $> p$

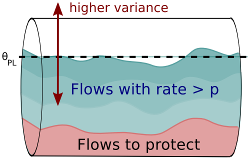

- Link reach saturation due to high variance of traffic

A Poisson approximation (2/2)

Flows of peak rate $> p$

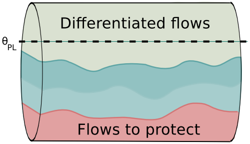

- Link reach saturation due to high variance of traffic
- Flows of higher peak rate backlogged
- PL decreases, new admissions

A Poisson approximation (2/2)

Flows of peak rate $> p$

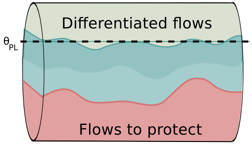

- Link reach saturation due to high variance of traffic
- Flows of higher peak rate backlogged
- PL decreases, new admissions

A Poisson approximation (2/2)

Flows of peak rate > p

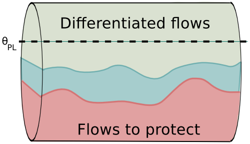

- Link reach saturation due to high variance of traffic
- Flows of higher peak rate backlogged
- PL decreases, new admissions

A Poisson approximation (2/2)

Flows of peak rate > p

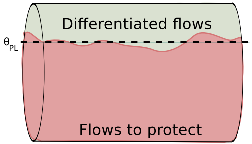

- Link reach saturation due to high variance of traffic
- Flows of higher peak rate backlogged
- PL decreases, new admissions
- New admission condition on instantaneous fair rate

A Poisson approximation (2/2)

Flows of peak rate $> p$

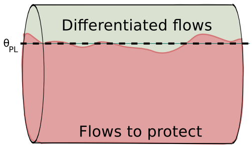

- Link reach saturation due to high variance of traffic
- Flows of higher peak rate backlogged
- PL decreases, new admissions
- New admission condition on instantaneous fair rate

Flows of peak rate $< p$

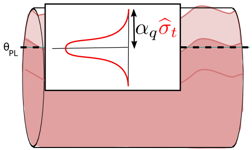

(Poisson) might be too conservative

$$\widehat{\sigma}_t^2 = min(A_t p, \sigma) \qquad \textbf{(MinVar)}$$

Slot size importance: $\tau = k\ L/p$

- $k < 1$: variance overestimated
- $k >> 1$: reactivity to load changes !

Outline

Evaluation: Simulation set-up

Topology

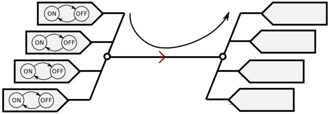

- $C = 10$Mb/s (depends only on C/p ratio)
- Packet size $L = 1000$ bytes
- Protected rate: $p = 100$ kb/s
- Sampling interval: $\tau = kL/p$, $k = 1, 2$

Traffic pattern

Poisson arrivals of on-off UDP flows with mean duration $T_h = 60$s

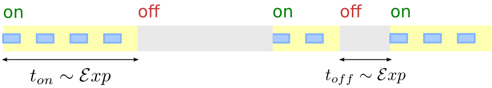

Flow peak rates

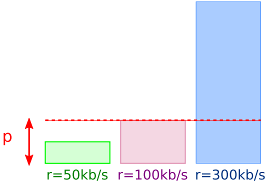

Evaluation: Utilization vs. overflow

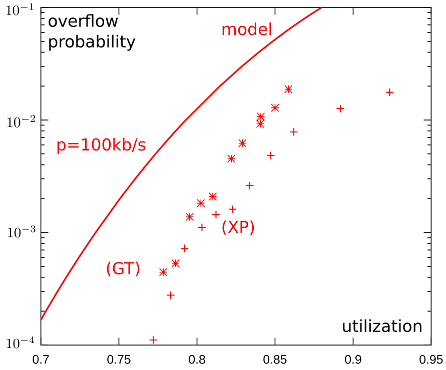

Evaluation: Utilization vs. overflow

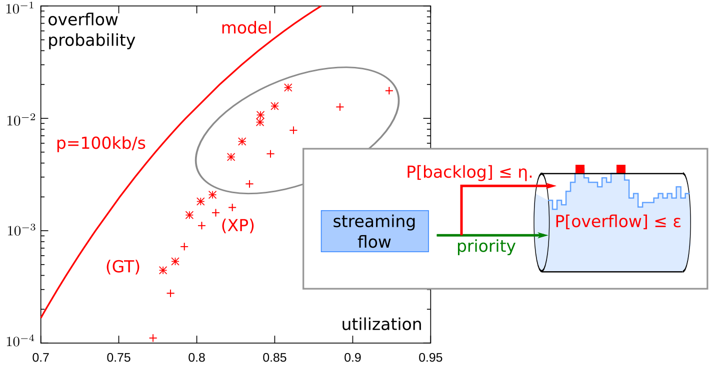

Evaluation: Utilization vs. overflow

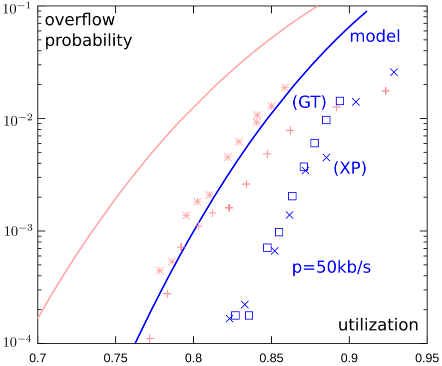

Evaluation: Utilization vs. overflow

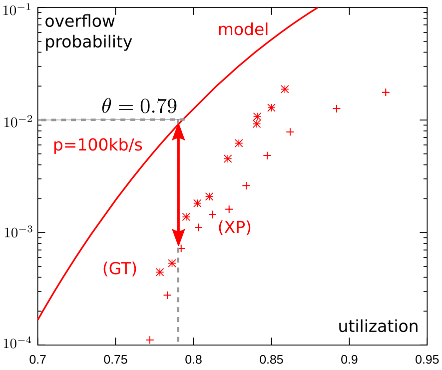

Evaluation: Utilization vs. overflow

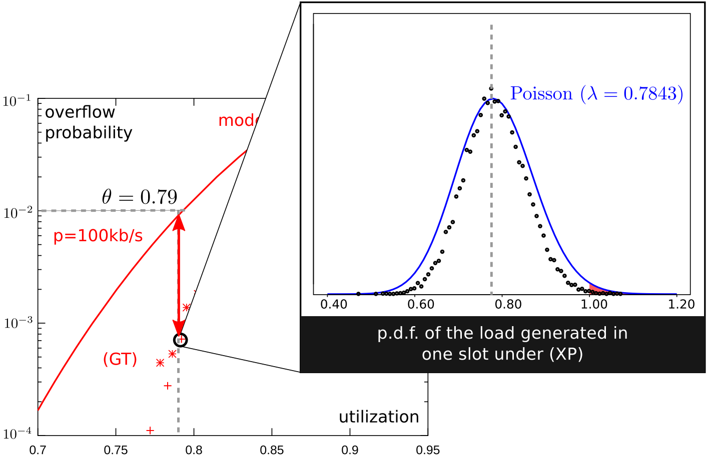

p.d.f. of the load generated in one slot under (XP)

Evaluation: Operating points

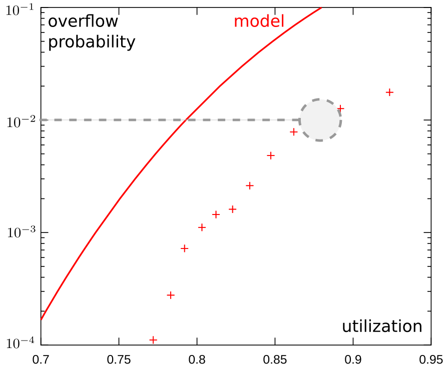

Evaluation: Operating points

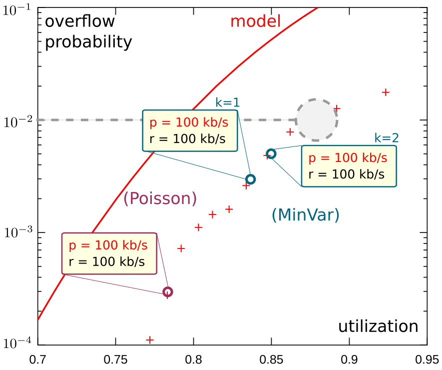

Evaluation: Operating points

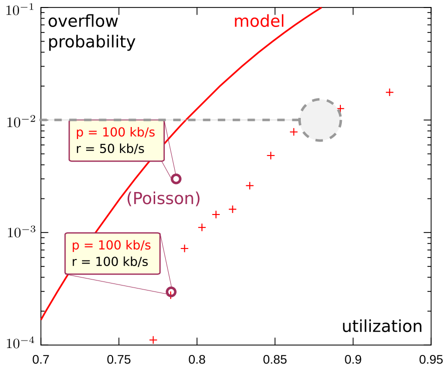

Evaluation: Operating points

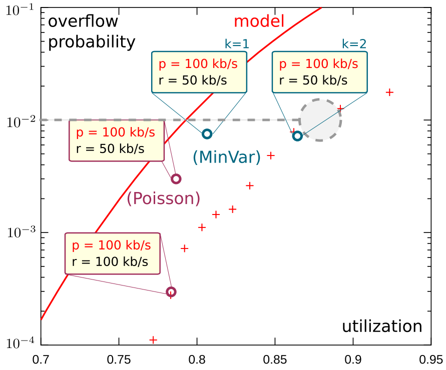

Evaluation: Operating points

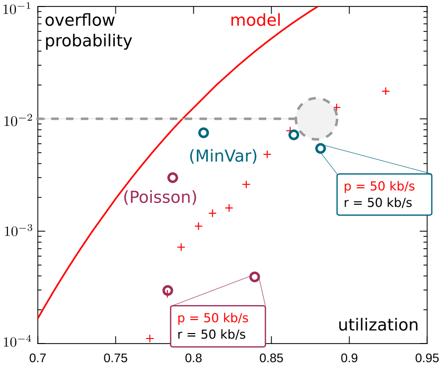

Evaluation: Differentiation

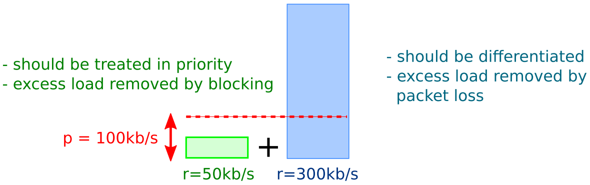

Evaluation: Differentiation

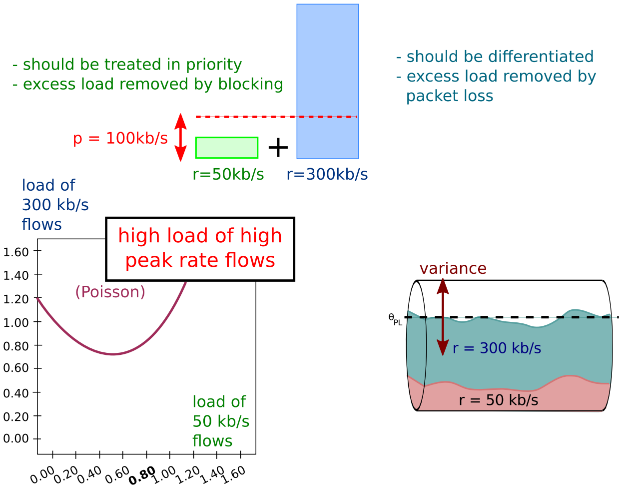

Evaluation: Differentiation

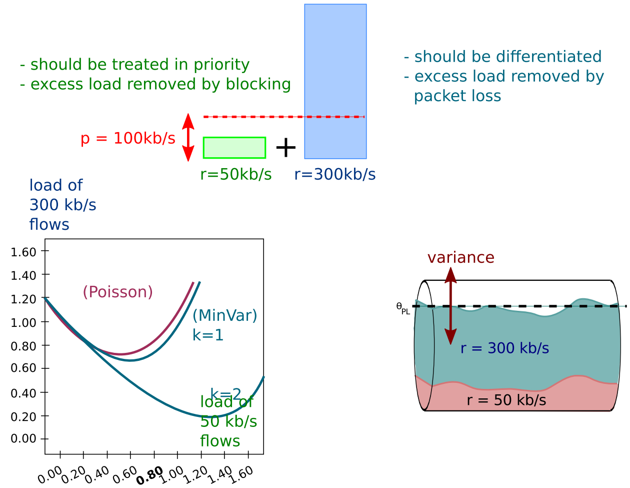

Evaluation: Performance under flashcrowd (1/2)

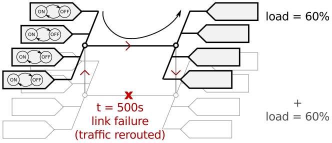

- We introduce a limit on the number of admission per slot
- No flow termination indication: hard to introduce back-off strategy

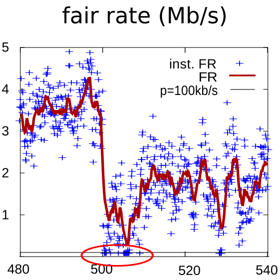

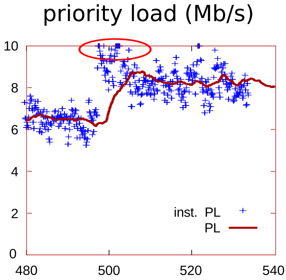

Evaluation: Performance under flashcrowd (2/2)

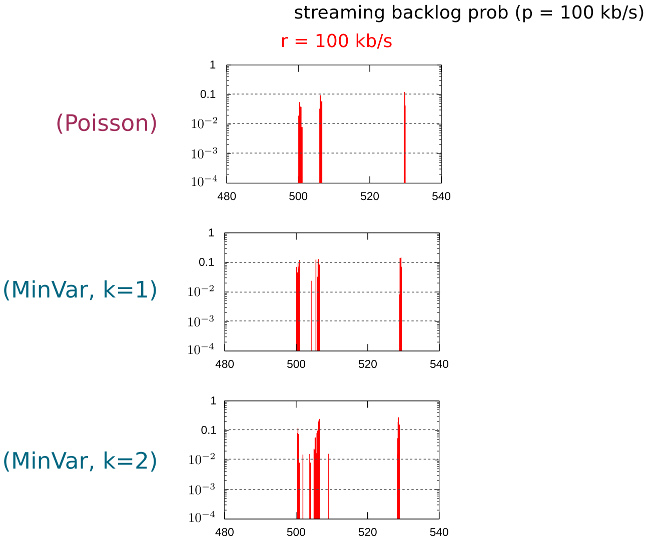

Evaluation: Performance under flashcrowd (2/2)

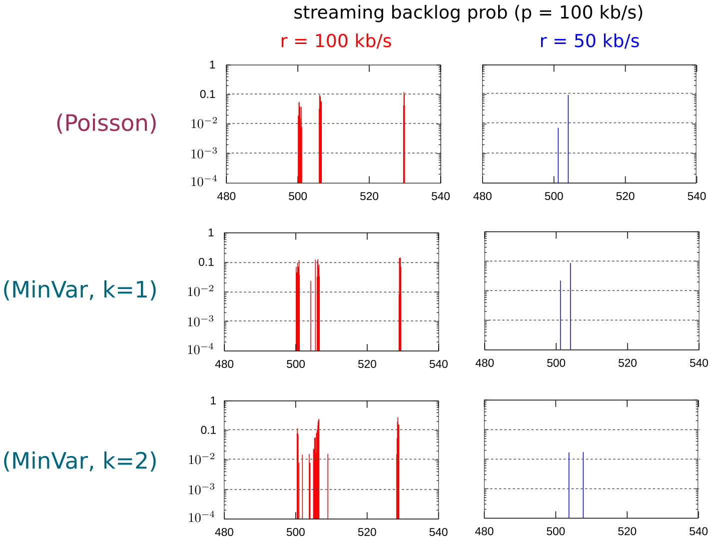

Outline

Conclusion

Use of Cross-Protect to offer performance guarantees to streaming and elastic traffic.

We have adapted a simple MBAC algorithm to protect streaming flows and allow for service differentiation.

Comprehensive set of simulation to demonstrate its performance.

Most problematic case is flashcrowd scenario

- (esp. heavy tail, TCP traffic)
- How to improve and react to the detection of such events ?
- Maybe a need to introduce flow preemption schemes. . .

Contact: jordan.auge@lip6.fr

Backup slides

Evaluation: Differentiation (details)

(Poisson)

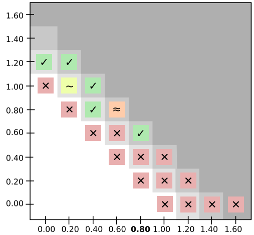

(MinVar, k=1)

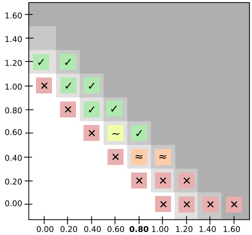

(MinVar, k=2)

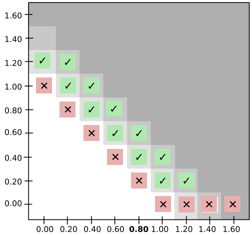

$p = 100\text{kb/s}$, 2 flow classes:

$r_2 = 300\text{kb/s}$

$r_1 = 50\text{kb/s}$

Level of differentiation: ✓ Full ~ Good (after some time) ≈ Episodic × None

Evaluation: Impact of jitter (1/2)

CBR flows typically acquire jitter in network routers

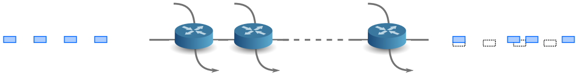

Simulation: Poisson stream of packets during on period [Better than Poisson conjecture]

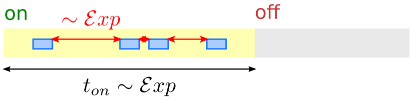

Evaluation: Impact of jitter (2/2)

Packet bursts are served at the fair rate

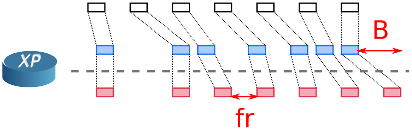

- No loss in simulation: packets are delayed instead of being dropped
- Unjittering as a supplementary advantage of Cross-Protect (provided sufficient buffer space)