An Admission Control based on a Knowledge Plane

Work-in-progress

RESO

March 23 2012

Motivation

- Growing traffic volume
  - e.g., video, P2P, video games
- Accept or Reject a new incoming flow ?
  - too many acceptances => performance collapse
  - too many rejections => low level of resource utilization
- State-of-the-art Solutions
  - measurement-based approach
  - implemented on each link
  - how to calibrate their parameters?
    - too specific to scenarios ☹

Our approach

- A communication link <=> a queue with a single server
- Yes, but which queue to choose?
  - *G/G/1* queue, but how to solve it?
  - *M/G/1* queue, simple solution for its steady-state but general enough?

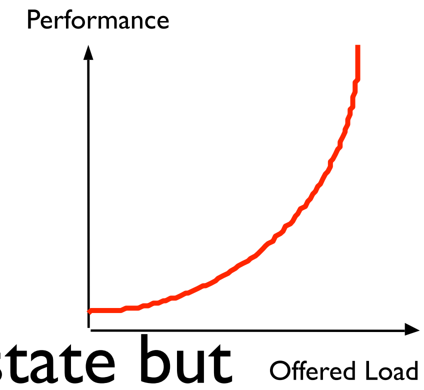

- How to choose the queue parameter values?
  - in a dynamic and automatic way
  - need for a procedure to set them from measurements collected on the link

Major Steps

- Measurements
  - on short timescale, collect couples of <*Xi*,*Pi*> with
  - *Xi* = average throughput,
  - *Pi* = packet loss rate or queueing delay
- Modeling the behavior
  - too many data => cluster <*Xi*,*Pi*> into k centroids
  - apply *High-Level Modeling* to find the fitting *M/G/1* queue
  - i.e., find f_{MG1} such that $Pj \simeq f_{MG1}(Xj) \quad \forall j$
- Decision algorithm
  - given the rate *r* of the incoming flow, workload growth projection
  - i.e., $f_{MG1}(Xi+r) < P^*$ with P^* = max. tolerable value

Illustration

- Example
  - here, *Pi* = queueing delay in the buffer
  - Best *M/G/1* queue fit: gives μ=5.01, cv=2.02 and off=0.08

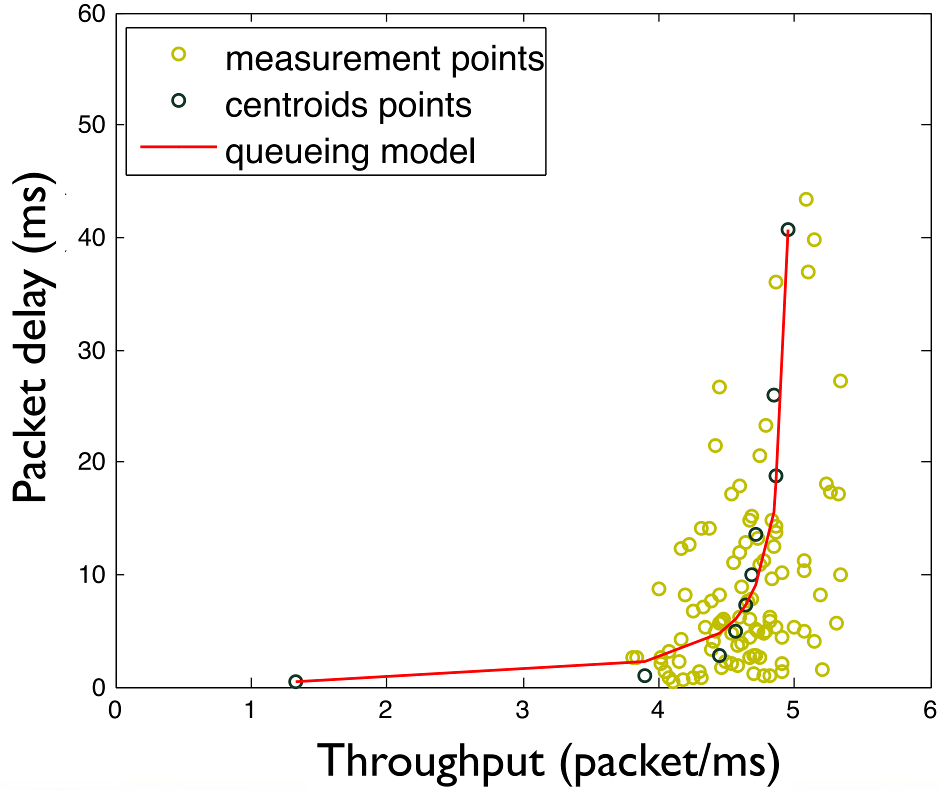

Shortcoming

- Too many measurement points <*Xi*,*Pi*>
  - Keep only 1000 points ☺
- But which of them?
  - the 1000 last points are the most recent ☺
  - but very likely to all look alike => loss of information ☹
  - instead, we split the possible range of throughput in 10 intervals,

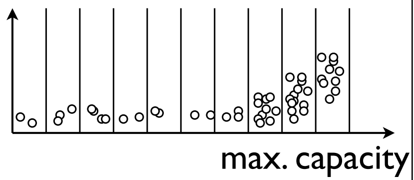

  - and we enforce the existence of 20 <*Xi*,*Pi*> in each of them
  - we guarantee the spatial diversity of centroids ☺

Illustration

- Example
  - here, *Pi* = queueing delay in the buffer
  - Best *M/G/1* queue fit: gives μ=5.01, cv=2.02 and off=0.08

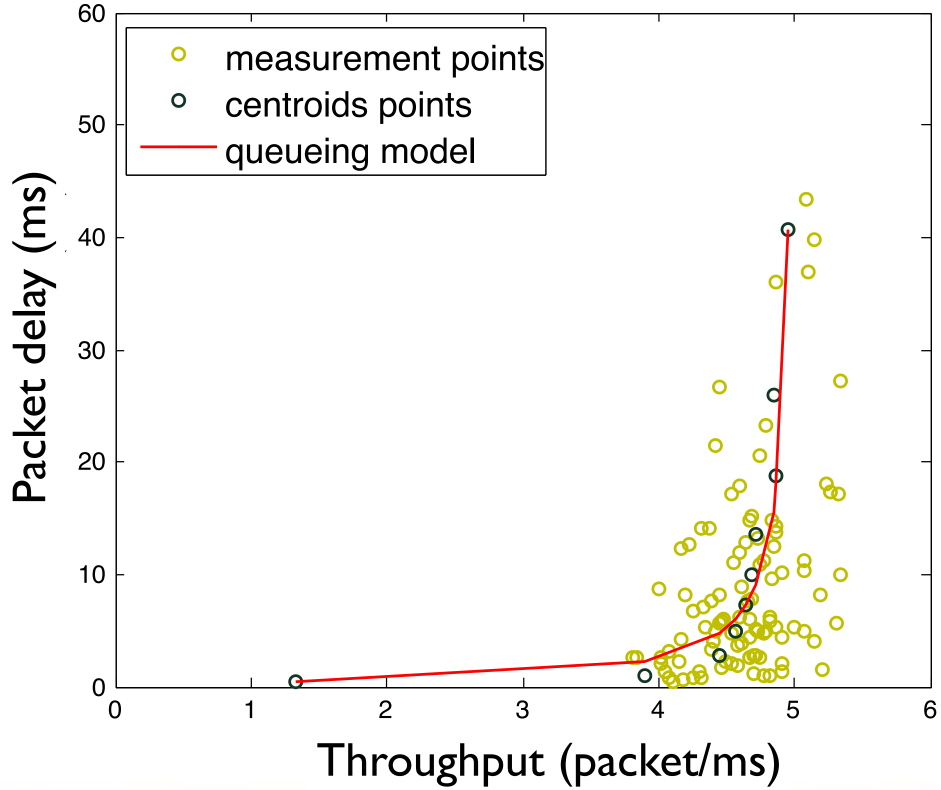

Video 1

Performance evaluation

- Link of 10 Mbps
- Maximum tolerable queueing delay, P*=10 ms
- Background traffic: real trace (~5.5 Mbps)
- Incoming flows :VBR flows (Poisson arrivals, mean duration 120s, mean rate 64 Kbps)

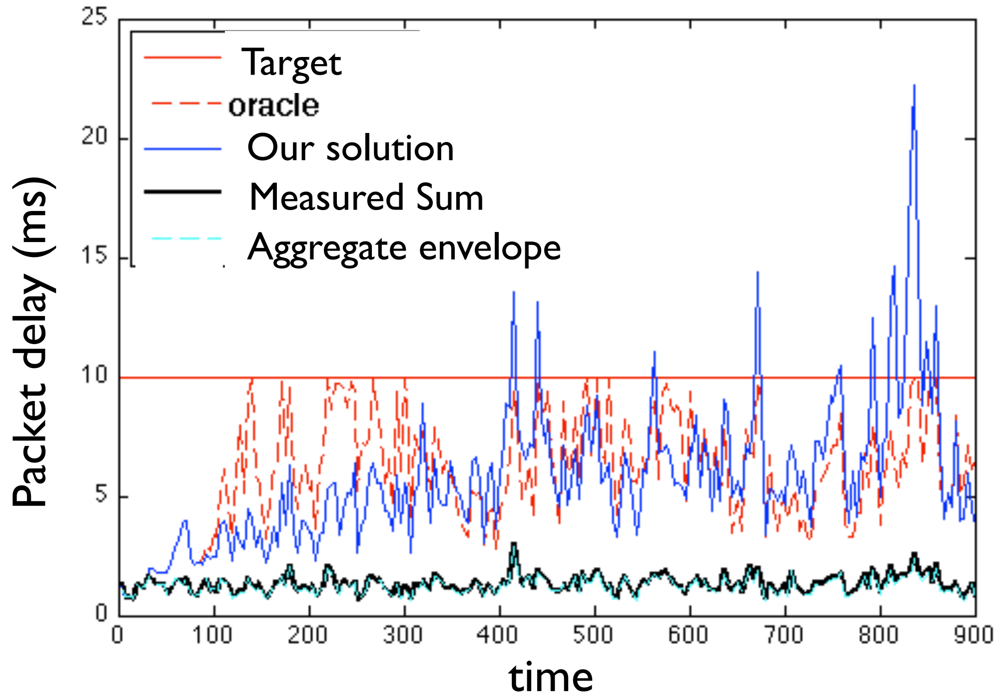

Video 2

Conclusion

- Data-driven and Evolutive solution
  - dynamic environment, no assumption on the traffic
- First step towards the definition of Knowledge Plane to sustain network management
- Semantic networking project of the common lab Alcatel Inria Bell Labs