EURESCOM SUMMIT
Heidelberg, October 2002

Flow aware Admission Control for a Commercially Viable Internet

T. Bonald, S. Oueslati, J. Roberts
France Telecom R&D

Requirements for the viability of a commercial network

- we assume
  - a competitive environment
  - viability must derive from the sale of transport services (no cross-subsidy)
- a commercial network is viable it federates a large customer base
  - high degree of sharing (individual demands << capacity)
  - assure adequate QoS economically - Technology is not the only consideration
- none of the currently proposed QoS models constitutes a cost-effective QoS solution
  - over-provisioning : simple and pragmatic engineering rule but not viable
  - Intserv, Diffserv : do not take sufficient account of the statistical nature of traffic

QoS and the statistical nature of demand

➔ assuring QoS relies on understanding the relationship between demand, capacity and performance

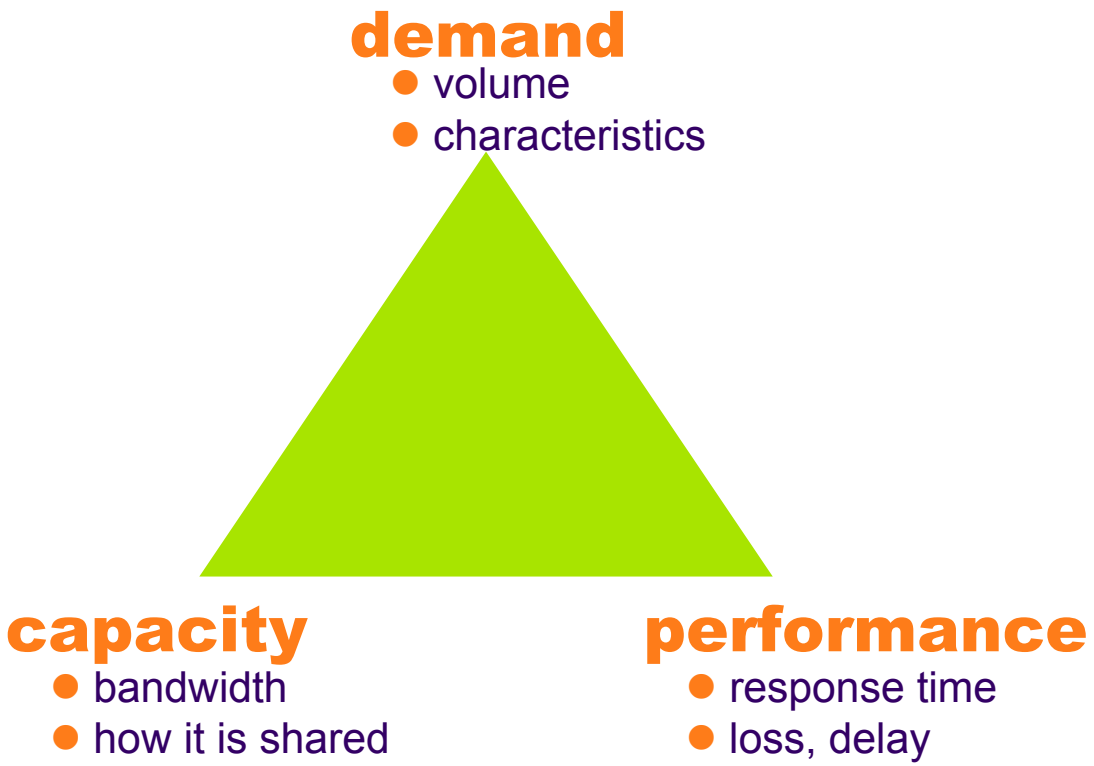

IP traffic variations

- traffic exhibits systematic day of week and time of day variations
- traffic is relatively stable in a wide busy period...
- ... with statistical variations about mean due to user activity

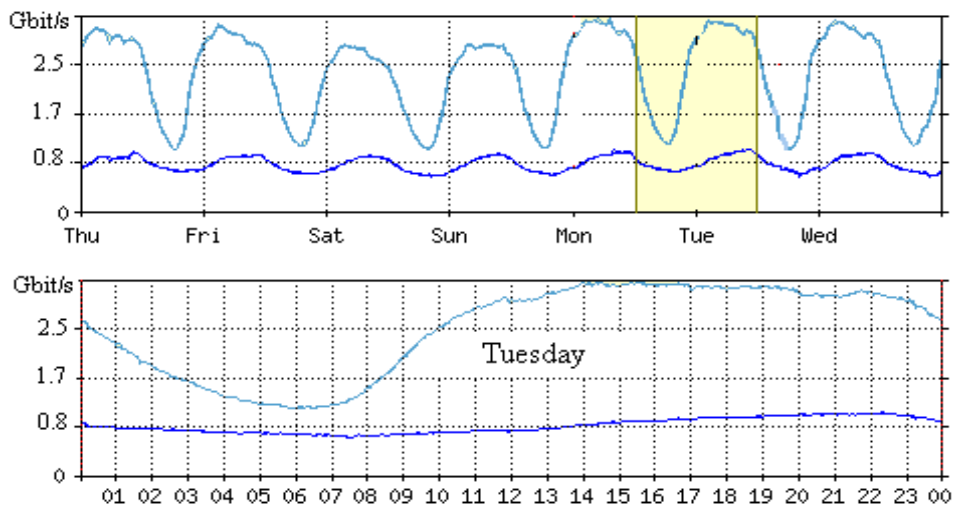

IP traffic variations

- traffic exhibits systematic day of week and time of day variations
- traffic is relatively stable in a wide busy period...
- ... with statistical variations about mean due to user activity
- a sizing objective:
  - enough capacity to meet QoS requirements (an "Erlang formula" for IP)...
  - ... with due allowance for overloads

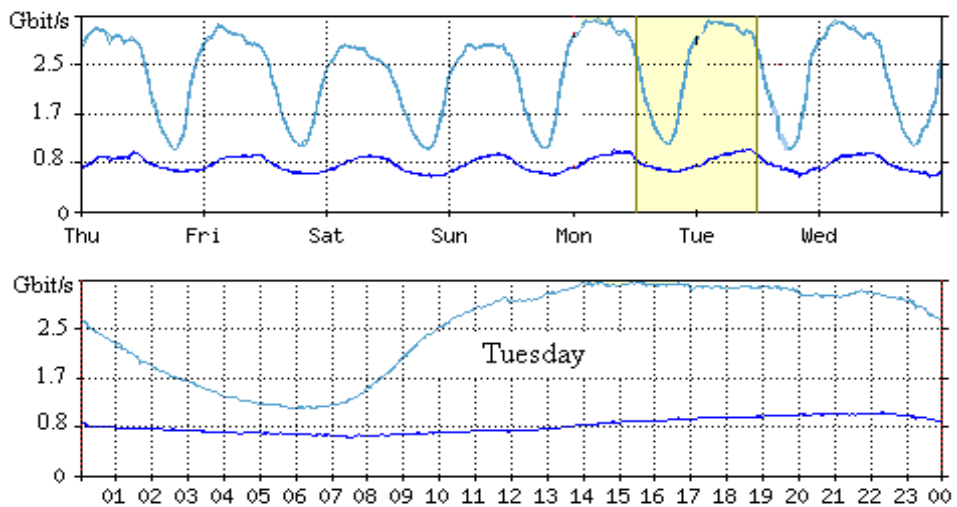

Composition of IP traffic

- by protocol:
  - TCP: 83% of packets, 91% of bytes
  - UDP: 14% of packets, 5% of bytes
- by application:
  - P2P, HTTP, ...
  - real audio, video, VoIP,...
- a simple robust classification:
  - elastic traffic (>90%)
  - streaming traffic (<10%)
- most traffic is (still) elastic
  - more than 90%

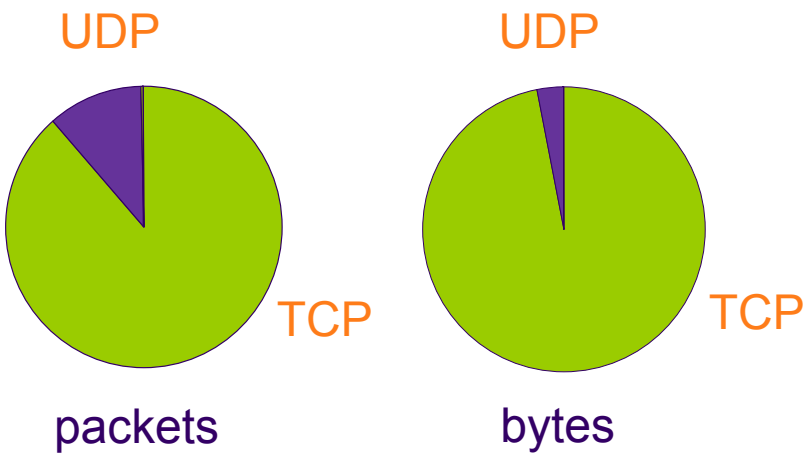

Packets, flows and sessions

- packets are part of "flows"...
  - a flow: all packets corresponding to one instance of a given application...
  - ... having the same identifier and occuring within a short time
- ... flows are part of "sessions"
  - a succession of flows and "think times"
  - relating to some homogeneous user activity (e-commerce, mail,...)

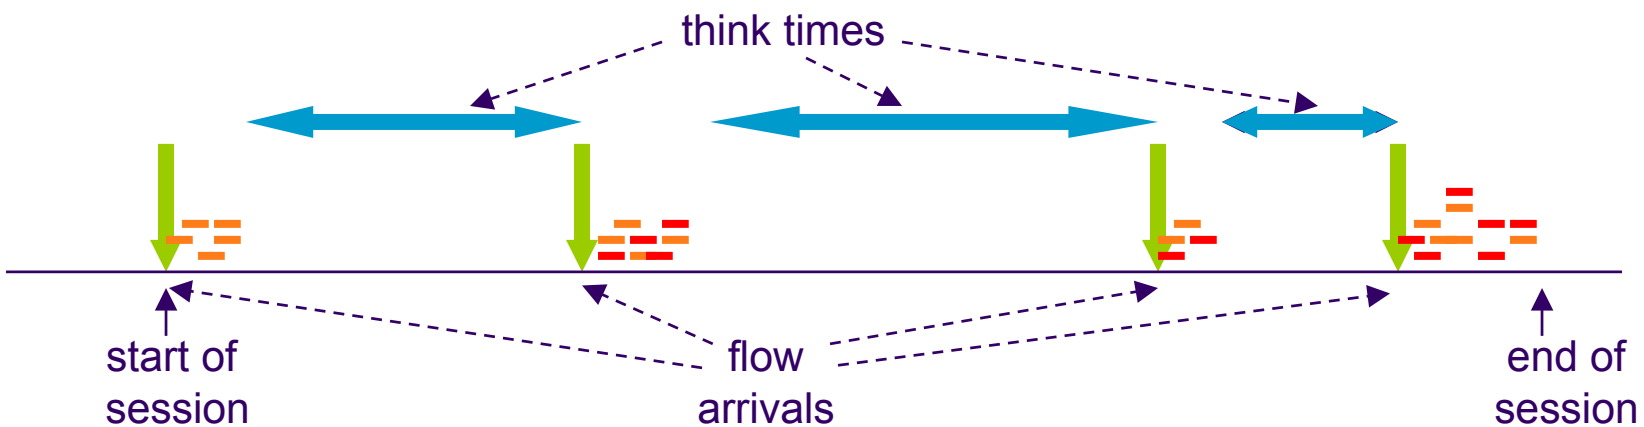

Packets, flows and sessions

- packets are part of "flows"...
  - a flow: all packets corresponding to one instance of a given application...
  - ... having the same identifier and occuring within a short time
- ... flows are part of "sessions"
  - a succession of flows and "think times"
  - relating to some homogeneous user activity (e-commerce, mail,...)
- modelling assumption: sessions occur as a Poisson process
  - in the busy period (like telephone calls!)

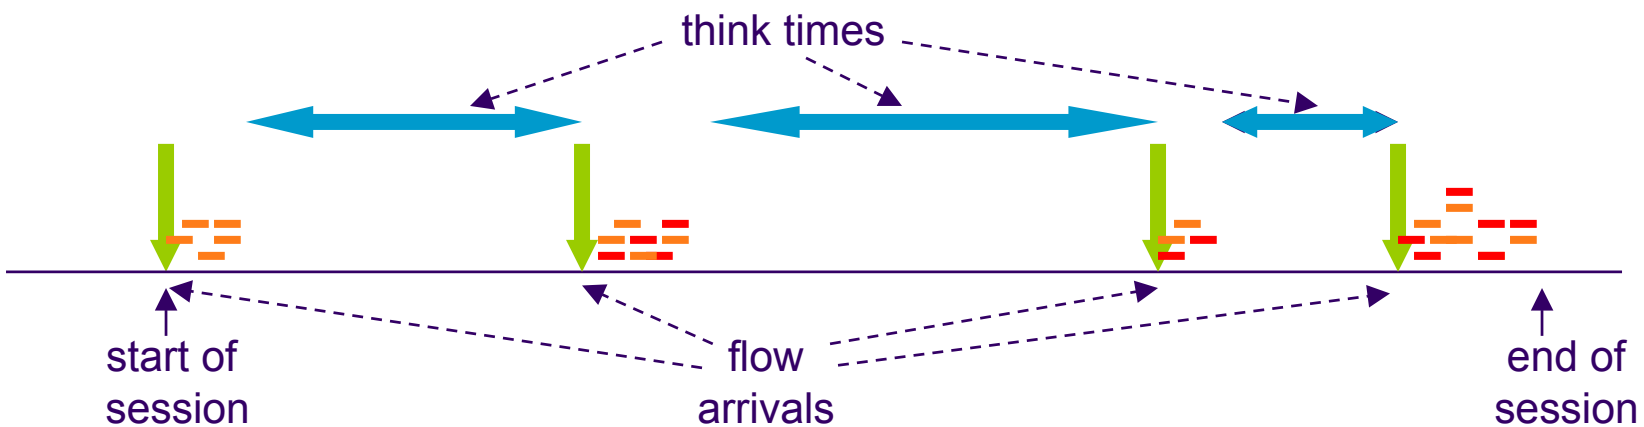

Performance model for elastic traffic

- TCP adjusts connection rates...
  - using the congestion avoidance algorithms
  - to (approximately fairly) share link bandwidth
- realized flow throughput depends on offered load
  - demand = arrival rate × average size (bit/s)
  - throughput ≈ capacity – demand (assuming fair sharing)

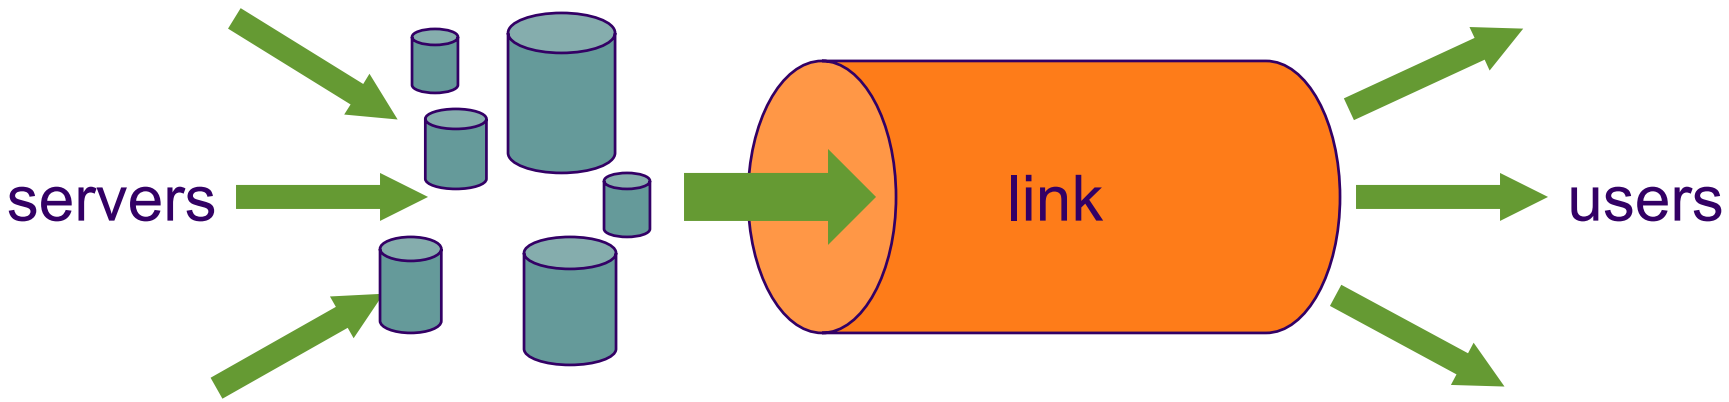

Simulation results: capacity 8 Mbit/s, offered load 90%

➔ visualization of per flow throughput

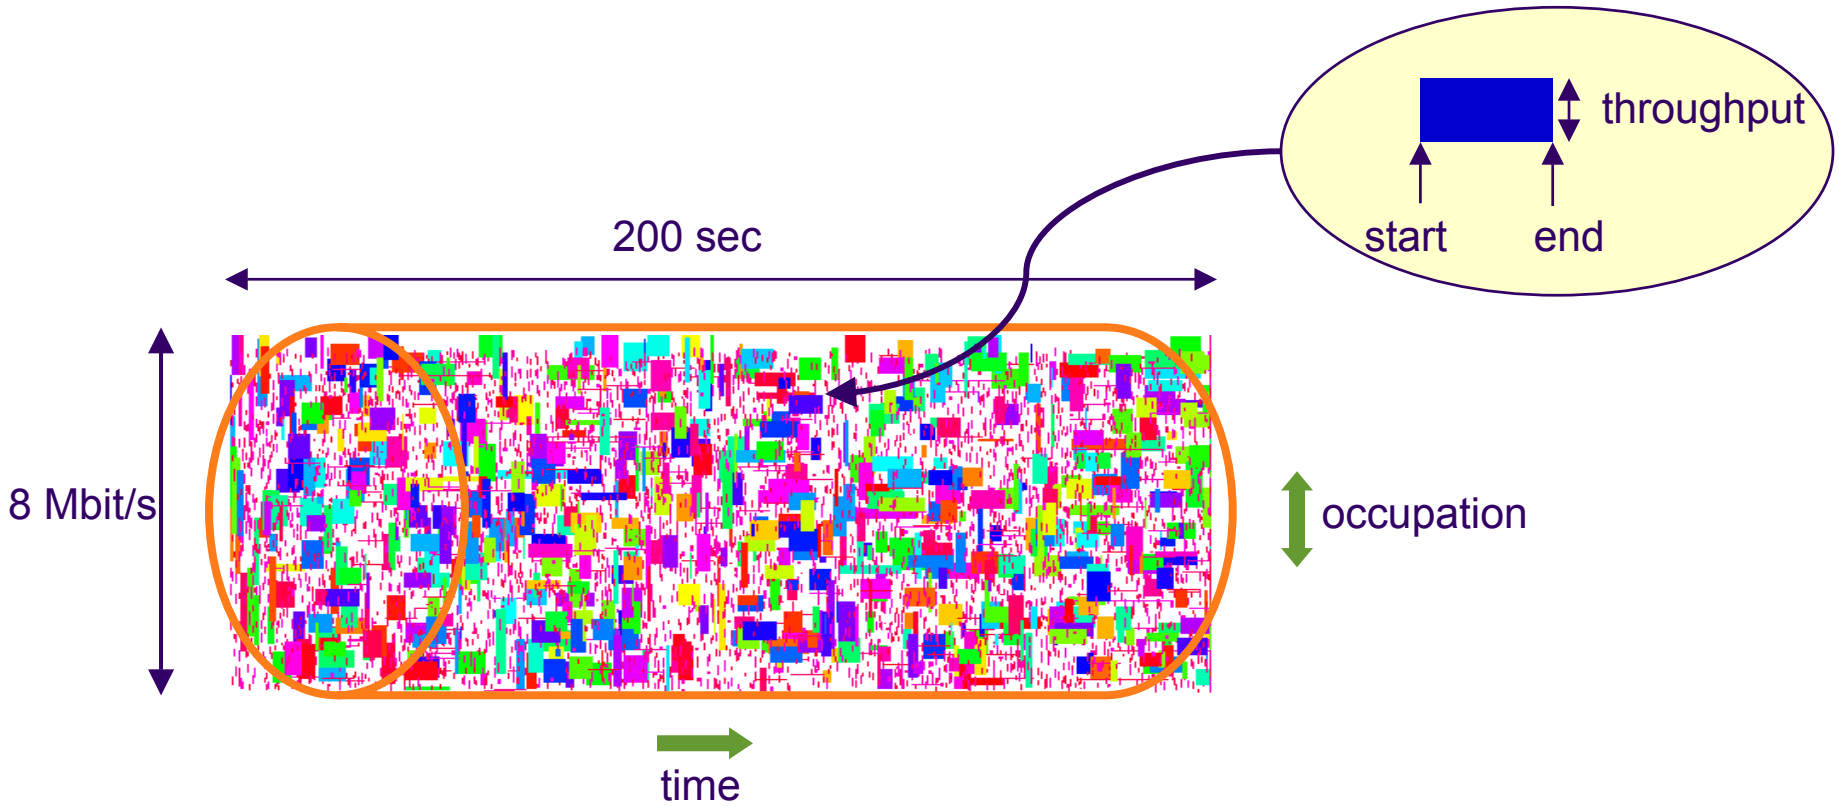

Simulation results: capacity 8 Mbit/s, offered load 90%

➔ visualization of per flow throughput

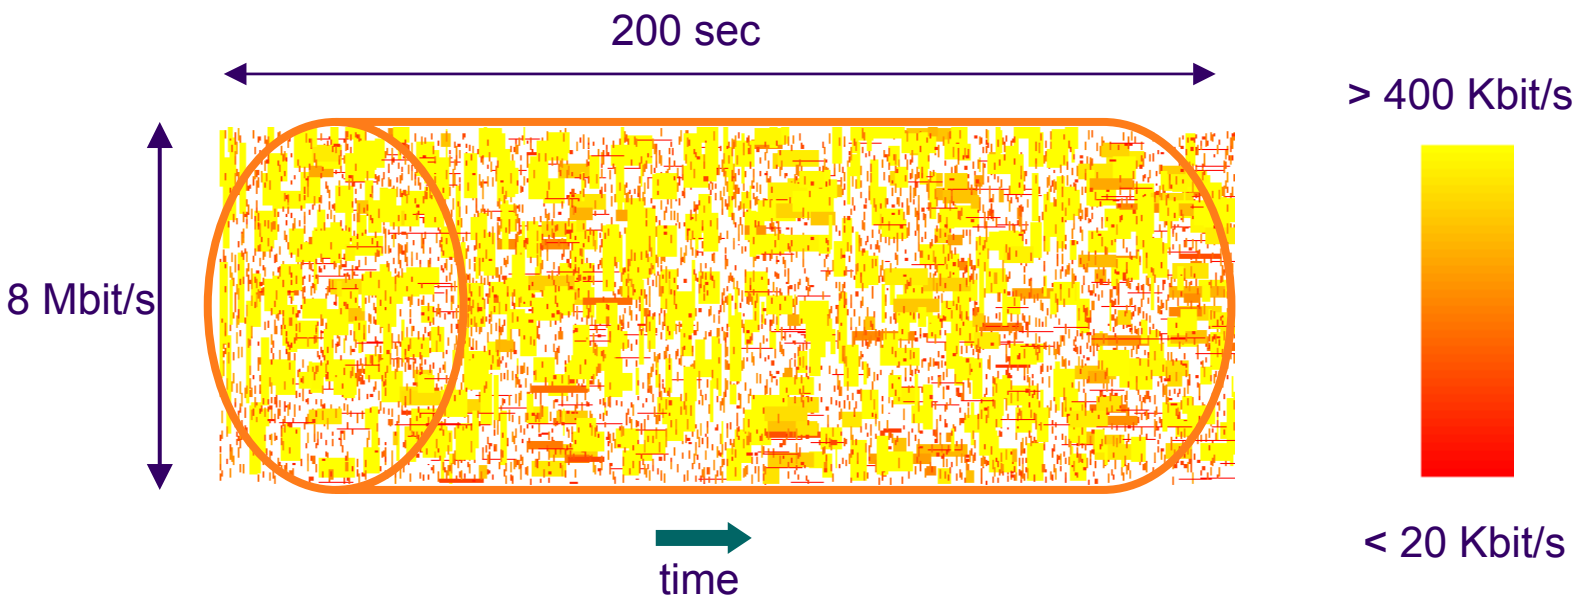

Simulation results: capacity 8 Mbit/s, offered load 90%

- visualization of per flow throughput
- the link is *transparent,* even at 90% utilization
  - throughput is determined by external limits
  - no scope for service differentiation

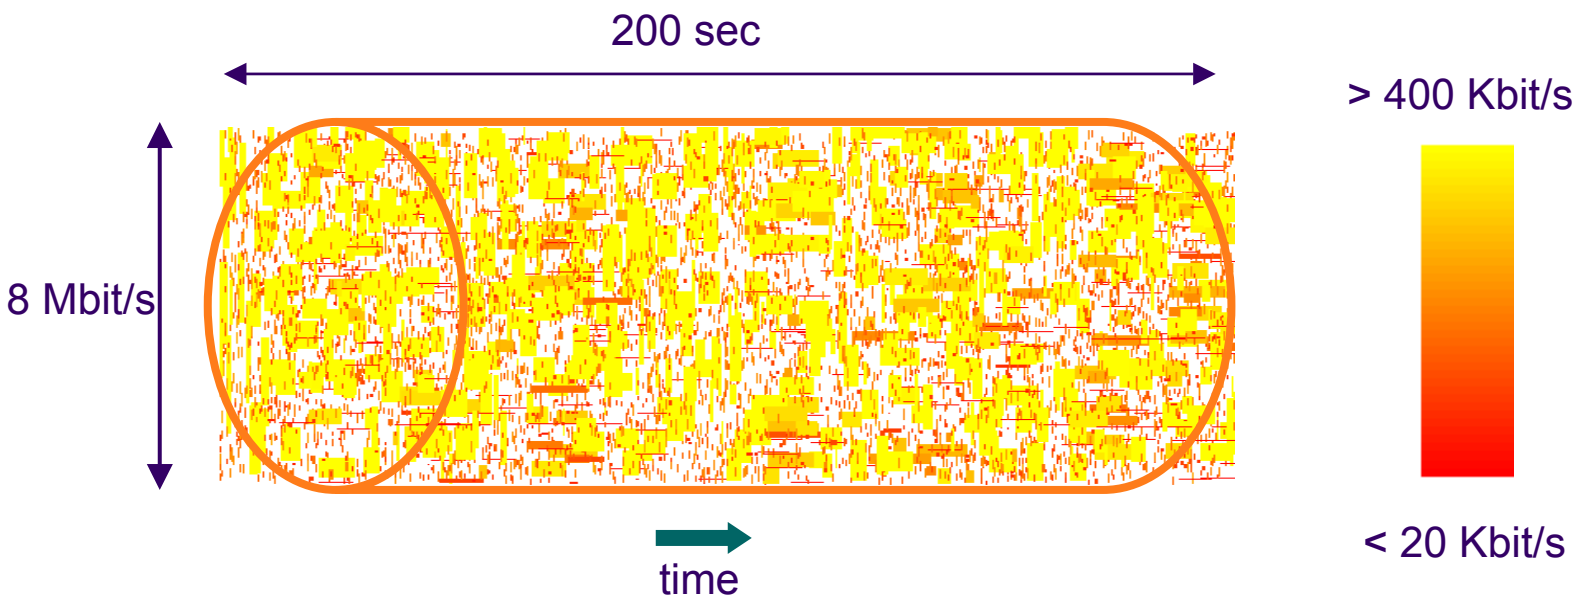

Simulation results: capacity 8 Mbit/s, offered load 90%

- visualization of per flow throughput
- the link is *transparent,* even at 90% utilization
  - throughput is determined by external limits
  - no scope for service differentiation
- transparency for streaming flows
  - give priority to streaming flow packets

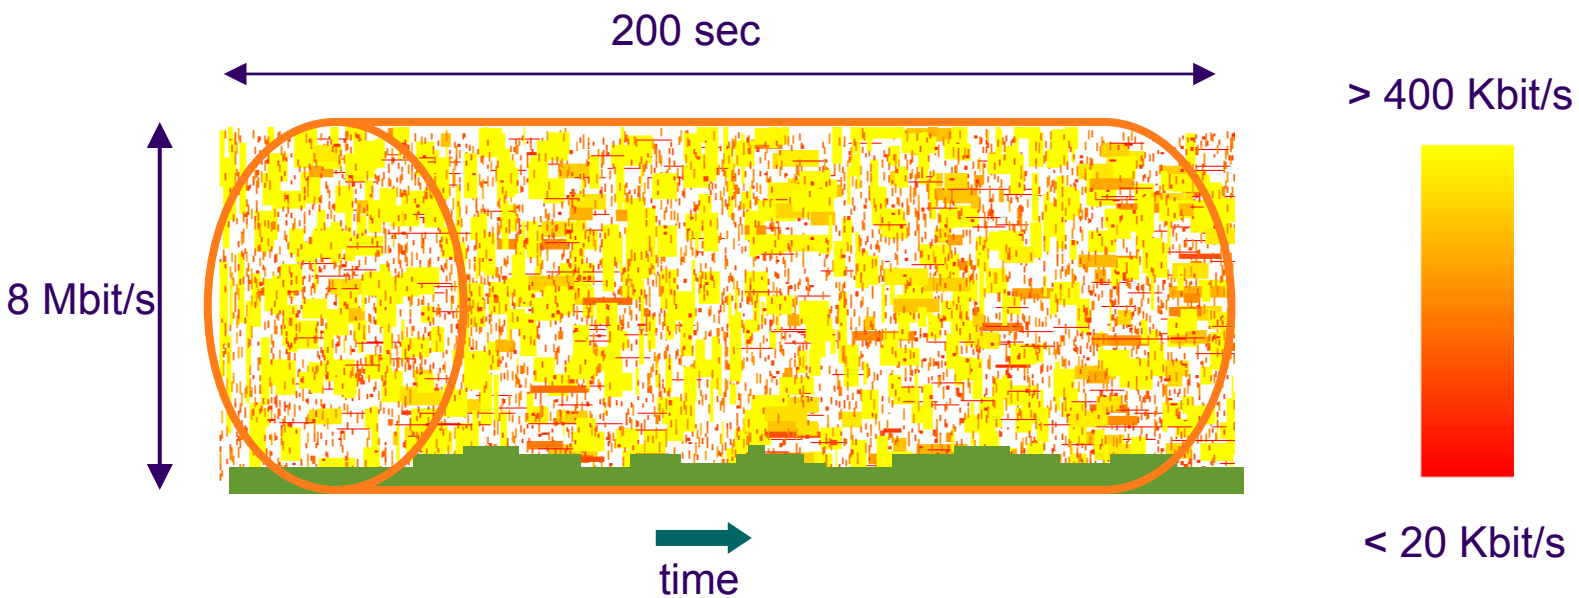

Impact of overload

- overloads occur (when demand > capacity)
  - failures
  - forecasting errors
  - traffic surges
  - deliberate choice (eg, some peering links)

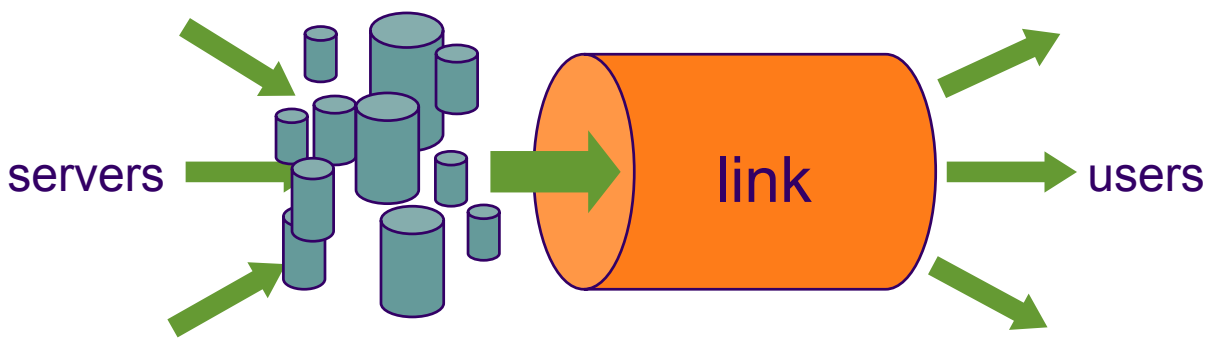

Impact of overload

- overloads occur (when demand > capacity)
  - failures
  - forecasting errors
  - traffic surges
  - deliberate choice (eg, some peering links)
- in overload the network is not transparent
  - arrival rate of new flows/sessions exceeds maximum service rate
  - throughput tends to zero ... until some flows are abandoned

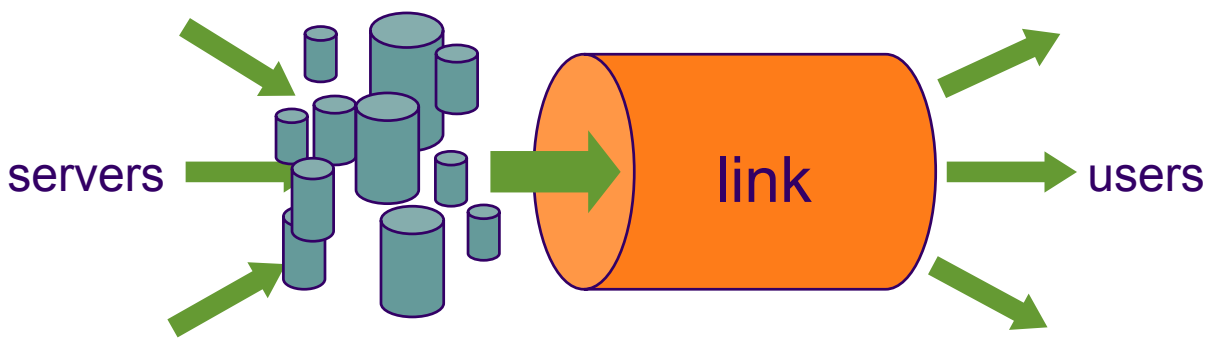

Simulation results: capacity 8 Mbit/s, load 140%

- per flow throughput tends to zero...
  - ... as new flows arrive faster than flows in progress complete

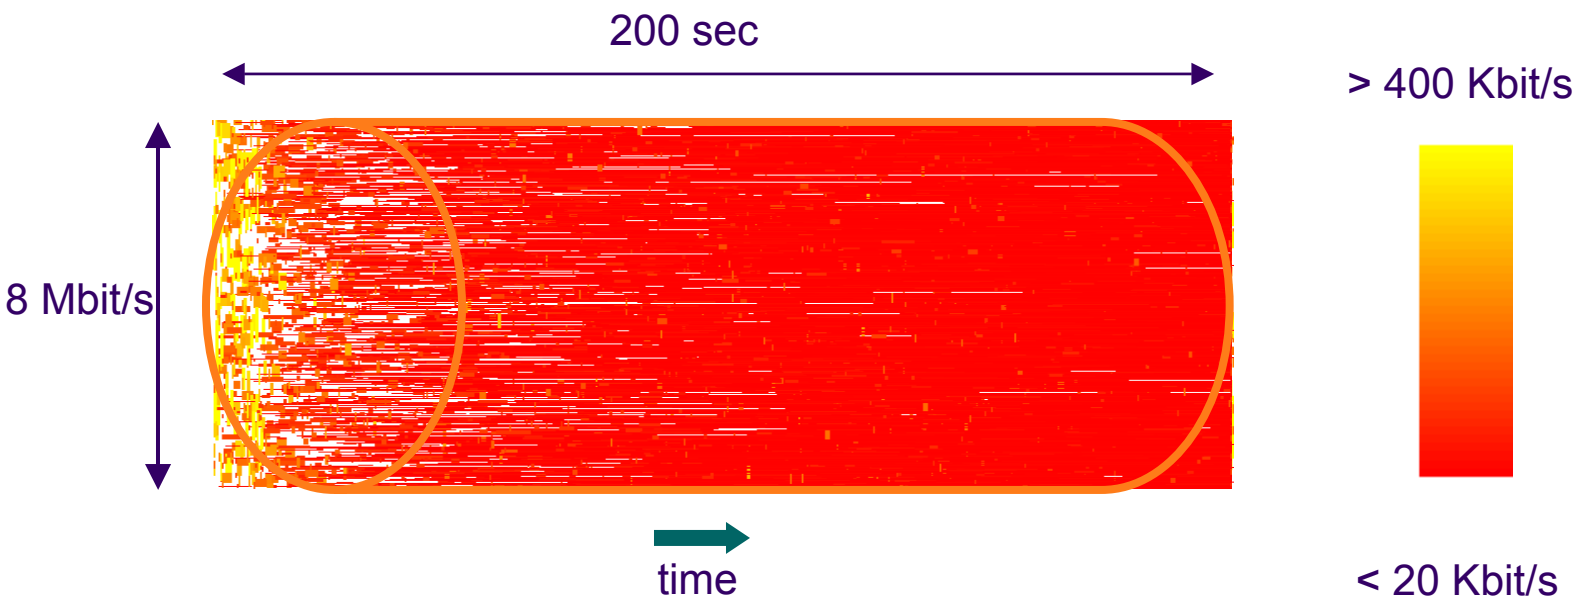

Service differentiation

- reserve sufficient capacity for premium traffic
  - service models: Intserv, Diffserv, MPLS
  - mechanisms: priority queues, WFQ,...

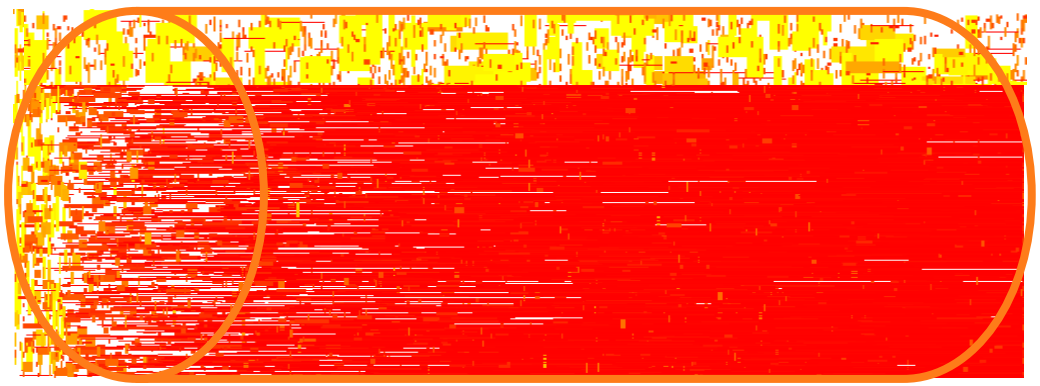

premium

best effort

Service differentiation

- reserve sufficient capacity for premium traffic
  - service models: Intserv, Diffserv, MPLS
  - mechanisms: priority queues, WFQ,...
- best effort traffic shares the rest
  - premium protected ... but most capacity is wasted

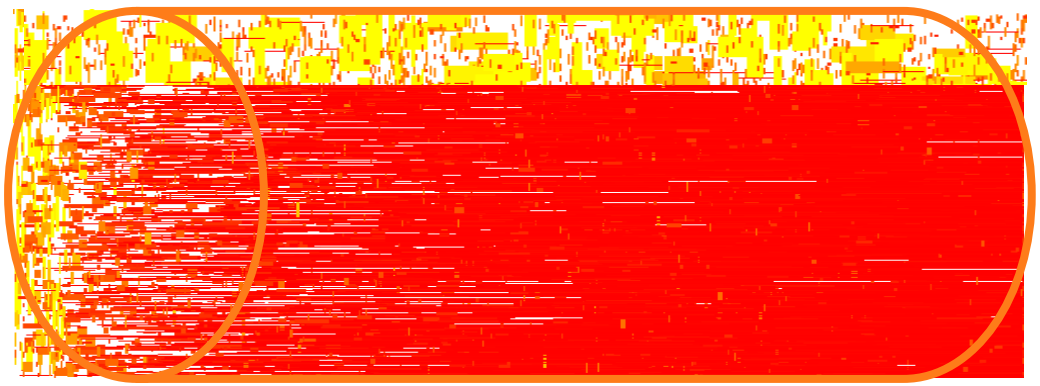

Service differentiation

- reserve sufficient capacity for premium traffic
  - service models: Intserv, Diffserv, MPLS
  - mechanisms: priority queues, WFQ,...
- best effort traffic shares the rest
  - premium protected ... but most capacity is wasted
- an alternative: preserve *transparency* by controlling *accessibility*
  - a flow-aware IP network architecture

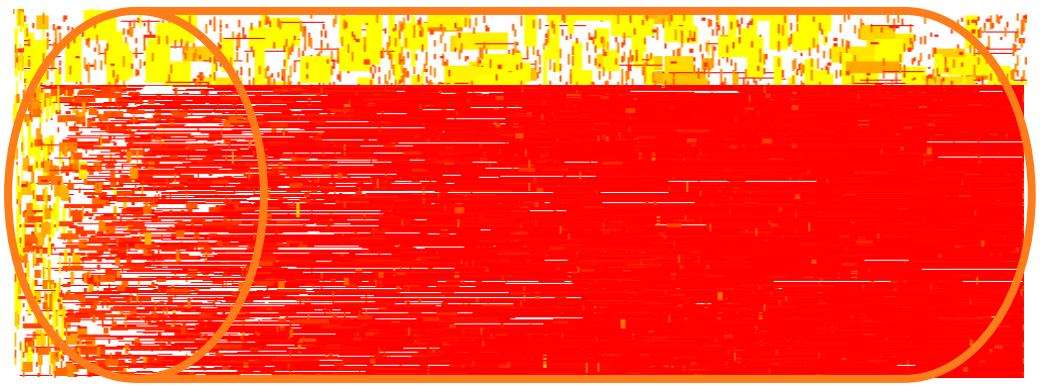

premium

best effort

Flow-aware admission control

- only accept new flows if their transparency is assured
  - reject new flows when their throughput would be too small
  - for both streaming and elastic flows

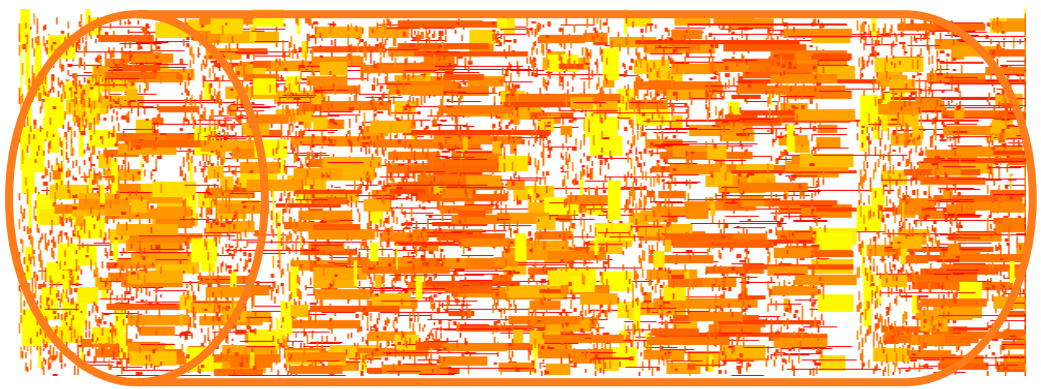

140% load

admission control

Flow-aware admission control

- only accept new flows if their transparency is assured
  - reject new flows when their throughput would be too small
  - for both streaming and elastic flows
- use different admission thresholds to provide differentiated accessibility
  - negligible blocking for premium traffic
  - all admitted flows are transparent and subject to charging

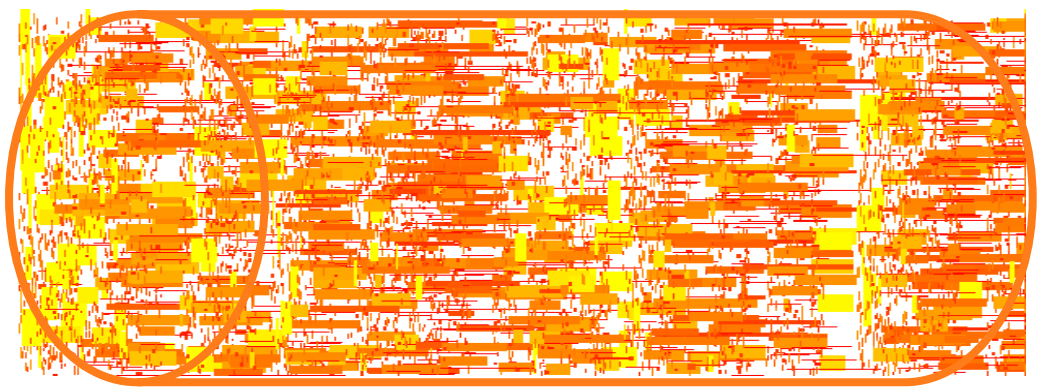

140% load

admission control

Implementing a flow-aware admission control

- no reservation, no signalling
  - an "enhanced best effort" architecture
  - localized admission control

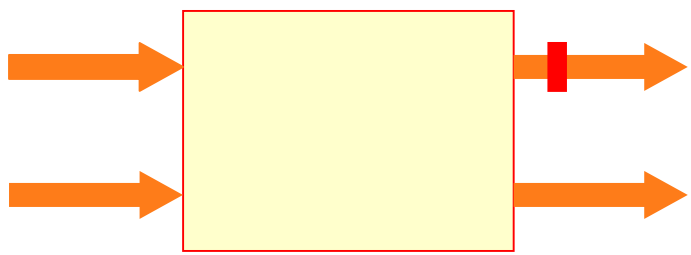

Implementing a flow-aware admission control

- no reservation, no signalling
  - an "enhanced best effort" architecture
  - localized admission control
- flow identification
  - user defined, using a flow label + IP addresses
  - a table of flows in progress

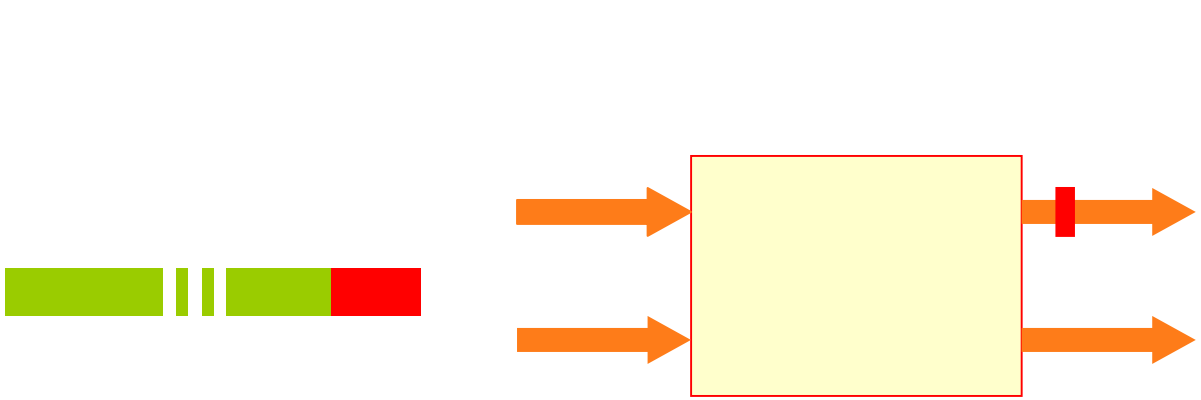

Implementing a flow-aware admission control

- no reservation, no signalling
  - an "enhanced best effort" architecture
  - localized admission control
- flow identification
  - user defined, using a flow label + IP addresses
  - a table of flows in progress
- implicit admission control
  - measure current level of congestion
  - reject new flows when necessary by packet discard

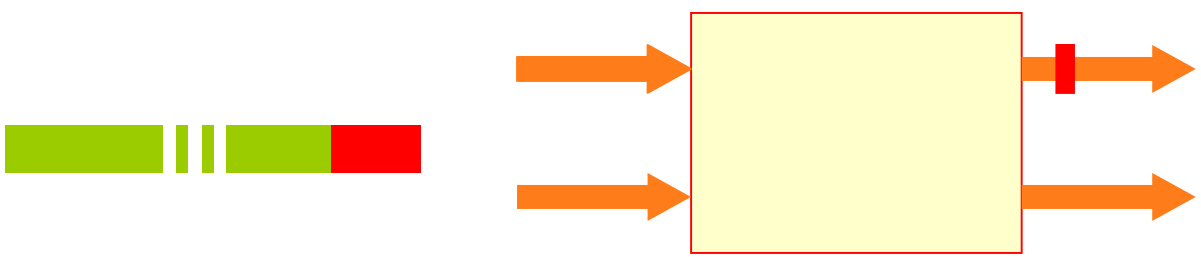

$flow_n$	$time_n$
$flow_m$	$time_m$

Conclusions

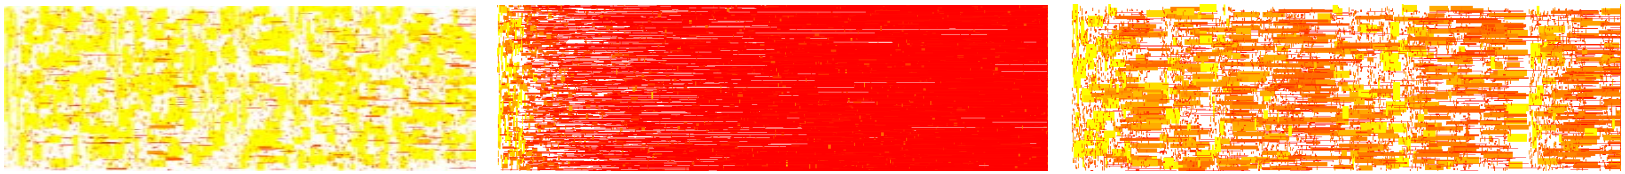

- in normal load (i.e., demand/capacity < 1-δ) the network is "transparent",
  - there is very little scope or need for service differentiation
- when demand > capacity, the network ceases to be transparent
  - (a phase change occurs)
- service differentiation can protect premium traffic...
  - but is wasteful and unnecessarily severe for best effort traffic
- flow by flow admission control constitutes an effective overload control
  - reject just the excess traffic, apply selective blocking
- realizing admission control: challenging but feasible
  - an implicit control: no signalling, no reservation!

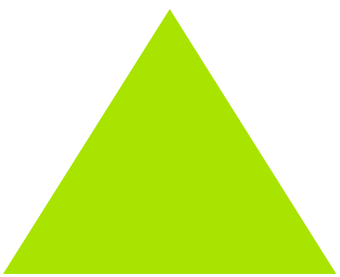

Conclusions

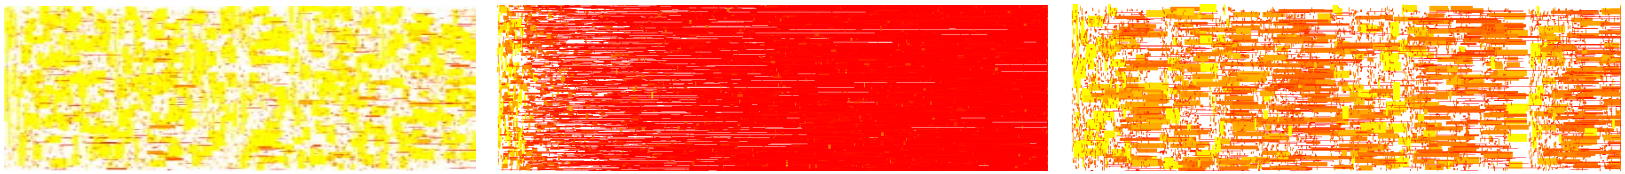

- in normal load (i.e., demand/capacity < 1-δ) the network is "transparent",
  - there is very little scope or need for service differentiation
- when demand > capacity, the network ceases to be transparent
  - (a phase change occurs)
- service differentiation can protect premium traffic...
  - but is wasteful and unnecessarily severe for best effort traffic
- flow by flow admission control constitutes an effective overload control
  - reject just the excess traffic, apply selective blocking
- realizing admission control: challenging but feasible
  - an implicit control: no signalling, no reservation!