NSF-COST Workshop
Hania, 23-25 June 2003

IP traffic and QoS control : the need for flow aware networking

Jim Roberts
http://perso.rd.francetelecom.fr/roberts
France Telecom R&D

QoS and the statistical nature of demand

➔ assuring QoS relies on understanding the traffic performance relationship

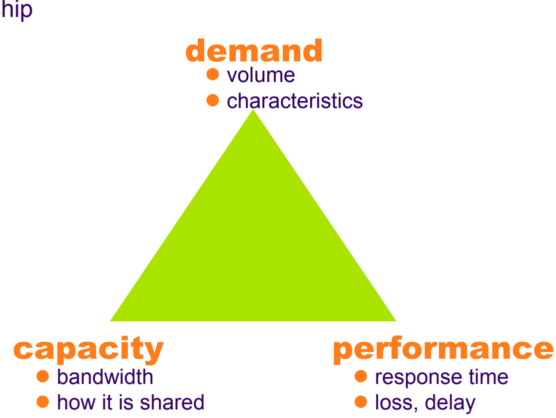

IP traffic is variable rate

- IP traffic is "self-similar"...
  - variability at all time scales
  - data, video, telephone,...
- ... because of
  - heavy-tailed size distributions
  - (bandwidth sharing by TCP)
- consequences
  - the failure of Poisson modelling?
  - the failure of the token bucket Tspec!

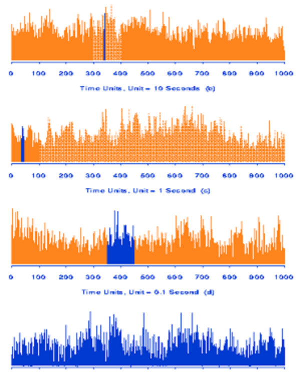

Ethernet traffic, Bellcore 1989

Prefer characterization at flow/session level

- 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!)

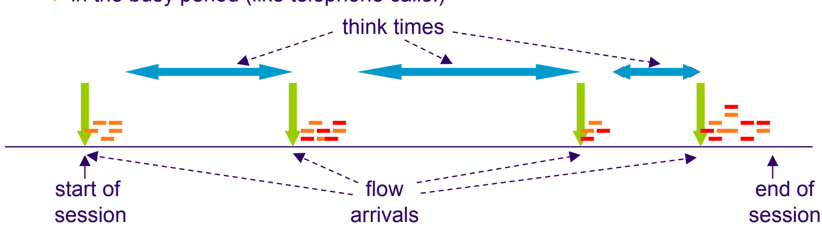

Traffic classification: streaming & elastic flows

- streaming traffic
  - real time audio and video applications
  - need signal conservation, i.e., "negligible" packet loss and delay
- elastic traffic
  - document transfers (as fast as possible): files, mail, Web, p2p,...
  - need for throughput conservation , i.e., "negligible" rate reduction with respect to external limits (access line, server capacity,...)
- this classification is robust
  - even as new applications emerge

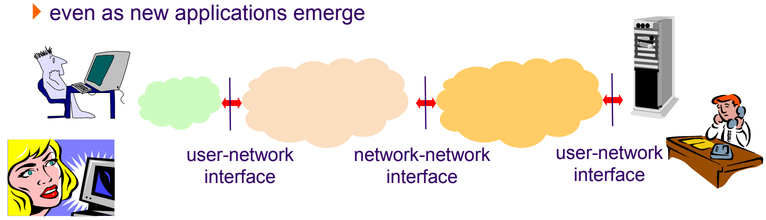

Performance of elastic traffic: an isolated bottleneck

- processor sharing model of bandwidth sharing
  - assume instantaneous fair shares
  - $\Rightarrow$ Pr [n flows in progress] = $\rho^n(1-\rho)$
  - $\Rightarrow$ E [throughput] = $C(1-\rho)$
- performance is *insensitive* to traffic characteristics
  - only necessary traffic assumption is Poisson session arrivals!
- its all a question of underload and overload
  - generally $C(1-\rho) >>$ external rate limits $\Rightarrow$ throughput conservation
  - when $\rho>1$, processor sharing model is unstable $\Rightarrow$ very bad performance

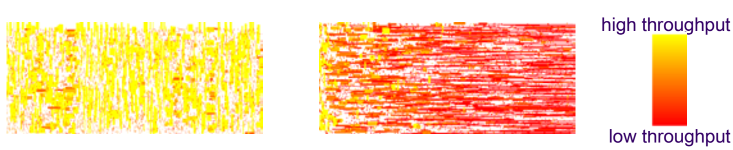

transparency ($\rho<1$) congestion ($\rho>1$)

Extension to networks

- notions of fairness
  - max-min fairness: ideal fairness (?)
  - maximized "utility", eg, proportional balanced, but what is utility in random traffic?
  - balanced fairness: an allocation that preserves processor sharing insensitivity (cf. papers by Bonald & Proutière - http://perso.rd.francetelecom.fr/bonald)
- how sensitive is real bandwidth sharing?
  - accounting for unfairness, TCP behaviour,...
  - still a question of underload and overload
    - $\rho < 1-\delta \Rightarrow$ transparency
    - $\rho > 1 \Rightarrow$ saturation

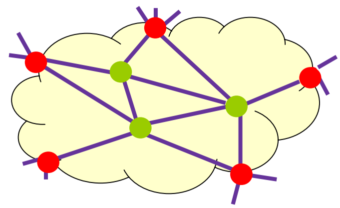

Integration of streaming and elastic traffic

- bufferless multiplexing for streaming traffic
  - ensures controlled loss, minimal delay
- fair sharing of residual capacity for elastic flows
  - integration realized by priority queuing and TCP
- performance model of an integrated link?
  - exact analysis is very difficult!
  - sensitive performance
  - quasi-stationary approximation
- a problem of local instability
  - when $A_e > C - \Lambda_s$
  - but not with admission control...

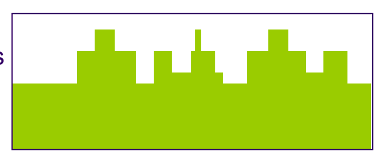

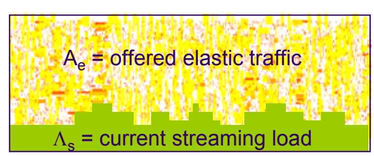

Packet level performance

- signal conservation means negligible packet loss and delay
- delay in one queue
  - $M/D_{MTU}/1$ as worst case limit of Σ D/D/1
- accumulation of jitter
  - the "NJ conjecture":
    - $M/D_{MTU}/1$ remains worst case in successive queues

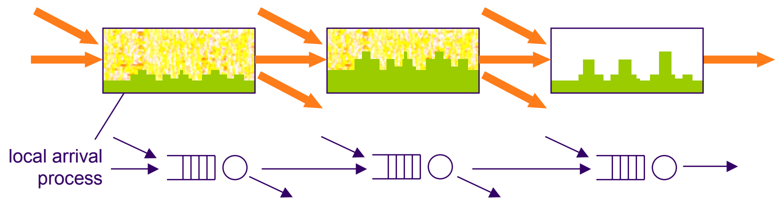

Traffic control

- admission control for streaming traffic
  - to ensure signal conservation
  - no *a priori* descriptors for variable rate flows
  - must use measurement based admission control (cf. Grossglauser &Tse, 2003)
- admission control for elastic traffic
  - to avoid congestion collapse (and ensure throughput conservation)
  - MBAC is easy!
  - using implicit control (on the fly flow identification, no signalling, no reservation)
- admission control on integrated link
  - facilitates MBAC for all flows

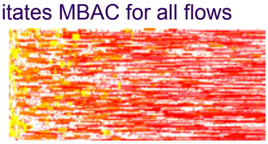

congestion (ρ>1)

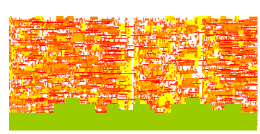

admission control (ρ>1)

"Cross protect": avoiding explicit differentiation

- ➔ signal conservation for flows of rate < fair rate
- ➔ throughput conservation by implicit admission control

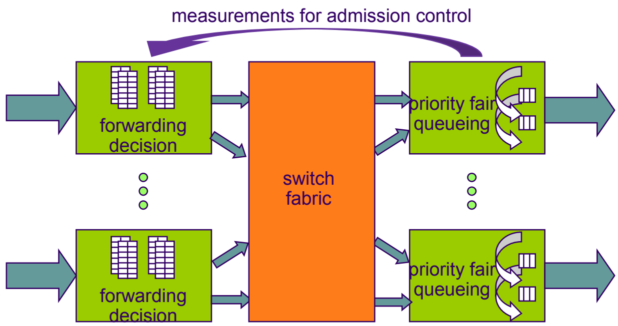

a cross protect router

Traffic performance in multiservice wireless networks?

- ➔ flows in wireless have a spatial traffic component
  - in addition to traffic characteristics relevant in wireline networks
  - wireless resource consumption depends on user position and mobility
- ➔ the resource is not just bandwidth!
  - users consume power and contribute interference
- ➔ little work accounting for random elastic traffic; see however:
  - S. Borst, *User-Level Performance of Channel-Aware Scheduling Algorithms in Wireless Data Networks*, Infocom 2003
  - T. Bonald & A. Proutière, *Wireless downlink data channels: User performance and cell dimensioning*, Mobicom 2003

Conclusion

- understanding the traffic performance relation
  - capacity – demand – performance
- characterize traffic at flow (and session) level
  - streaming and elastic flows
  - Poisson arrivals
- modelling elastic bandwidth sharing
  - processor sharing model and extensions
  - a question of underload and overload!
- integration of streaming and elastic traffic
  - difficult to model: sensitive and locally unstable
  - negligible jitter for streaming flow packets
- traffic control
  - measurement-based, per-flow, implicit *admission control*
  - facilitated in an integrated system
  - a new proposal: the *cross protect router*