What QoS for the future Internet?

Jim Roberts
INRIA

ITC 21
Paris, September 2009

QoS research and the future Internet

- a future Internet for enhanced security, mobility management and QoS
  - addressed diversely in research projects throughout the world
- major innovations are emerging like
  - content-centric networking
  - network virtualization
- but do we have a clear idea about enhancing QoS?
  - implement the models that have already been defined
  - or invent a new paradigm?

The dual role of QoS

- perform effective traffic management
  - to meet diverse application requirements
  - for delay, jitter, loss, throughput,...
- create a viable business model for the network operator
  - a range of differentiated services
  - to maximize revenue and avoid commoditization
- a source of some confusion!

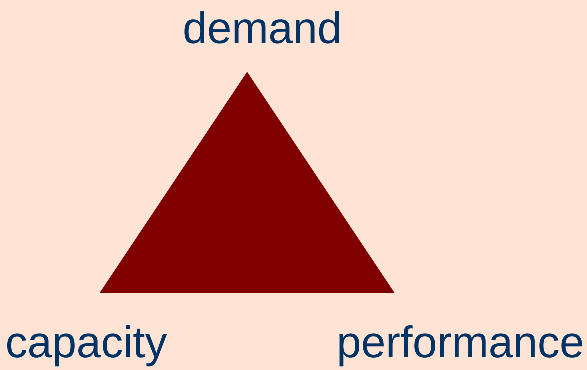

Outline

- Traffic and congestion
- Myths and mystification
- Managed or unmanaged Internet

Understanding traffic at flow level

- users experience quality at flow level
  - a flow is an instance of some application (document transfer, voice signal,...)
  - a set of packets with like header fields, local in space and time
- flows of four types
  - conversational, streaming, interactive data, background
  - with different requirements for latency, integrity, throughput
- an essential characteristic: the flow rate
  - constant or variable, high or low peak rate

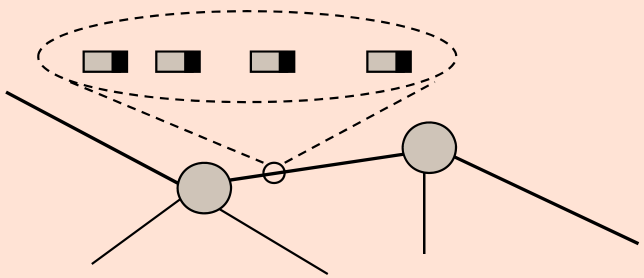

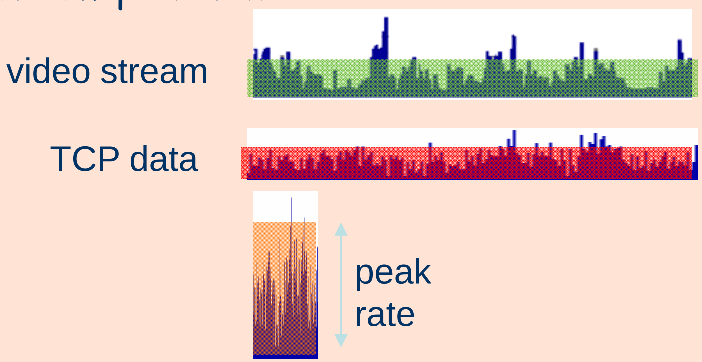

Two essential traffic characteristics

1. the mix of flow rates

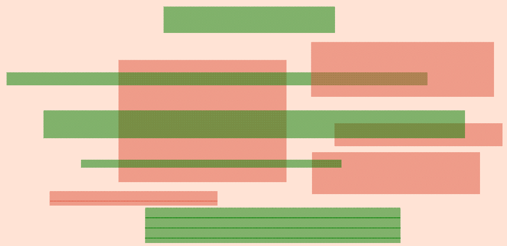

2. the link load

$$\rho = \frac{\text{flow arrival rate x mean flow size}}{\text{link capacity}}$$

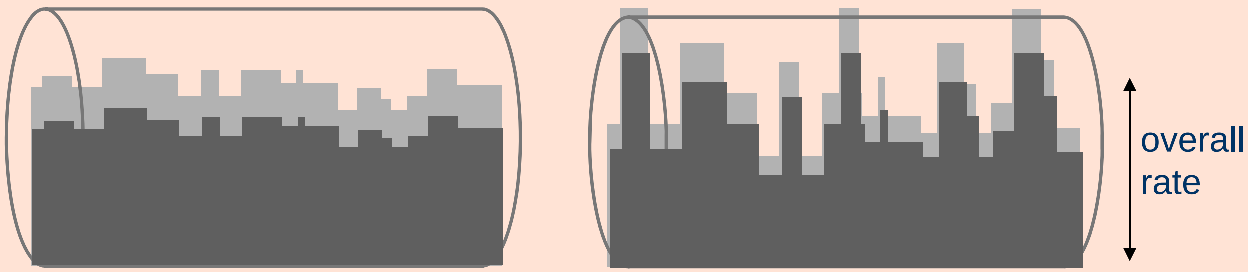

Three bandwidth sharing regimes

- transparent

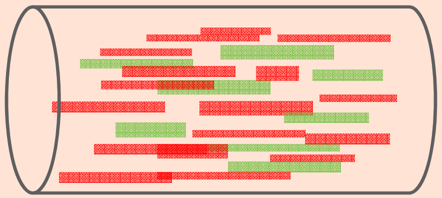

- elastic

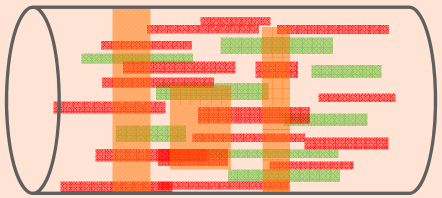

- overload

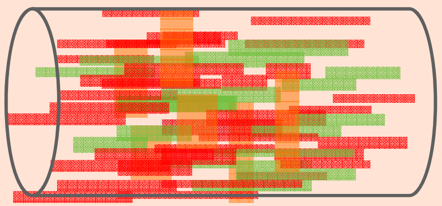

Transparent bandwidth sharing

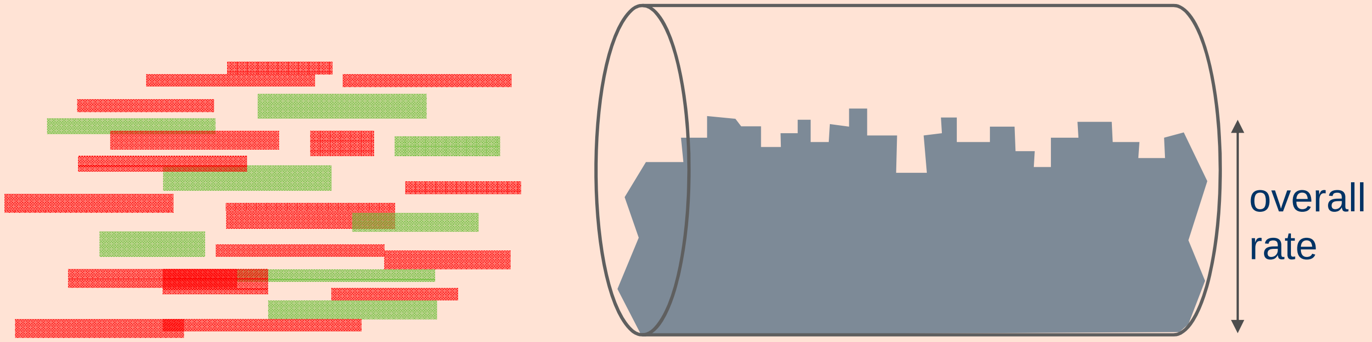

- all flows have relatively small rate
- offered load is somewhat less than capacity ⇒ no rate excess
- excellent quality for all
- this is "QoS by over provisioning"

Performance in the transparent regime

- predictable packet level performance
  - negligible delay at normal load
  - performance is insensitive to detailed traffic characteristics
  - not true if overall rate can exceed capacity!
- utilization can be high when peak rates are relatively small
  - not on access links
  - not when some flows have a very high rate

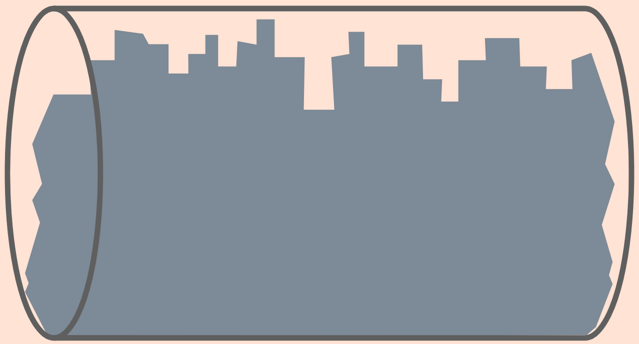

Elastic bandwidth sharing

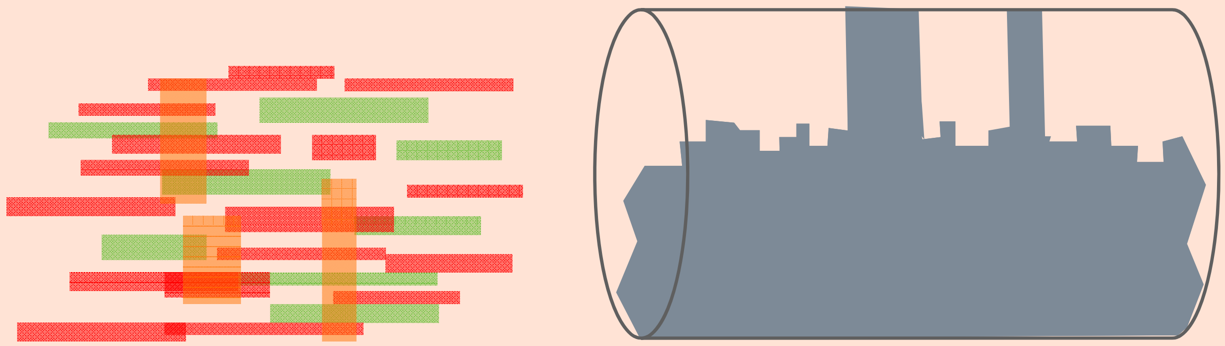

- some flows have high rate
- offered load somewhat less than capacity
- need to control sharing
  - to avoid loss to streaming flows
  - to ensure "fair" sharing

Performance in the elastic regime

- a “processor sharing” model of bandwidth sharing
  - assume all flows of unlimited rate, Poisson flow arrivals, instantaneous fair sharing
    - Pr [n flows] = $\rho^n (1 - \rho)$
    - E [throughput] = $C(1 - \rho)$
  - same results apply for more general traffic
- on high capacity links, most flows peak rate limited
    - E [throughput] $\approx$ min {peak rate, $C(1 - \rho)$}

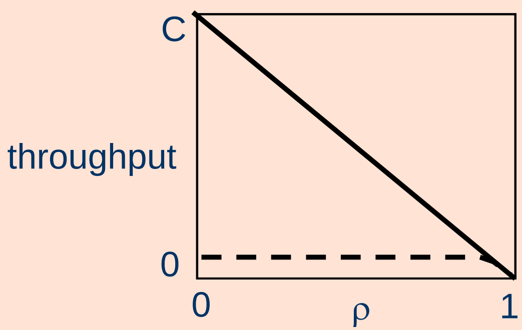

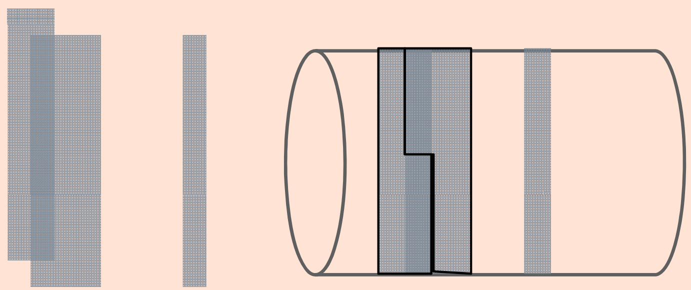

Controlling bandwidth sharing

- TCP realizes approximate fair sharing
  - but relies on end user cooperation
  - and new high speed TCPs are too aggressive
- apply economic incentives for congestion control?
  - eg, a "self-managed Internet" (F. Kelly, 2000)
  - but congestion pricing is unacceptable!
- or implement fair sharing in the network
  - difficult flow identification
  - but only rare flows are “elastic” in the core...

Overload and bandwidth sharing

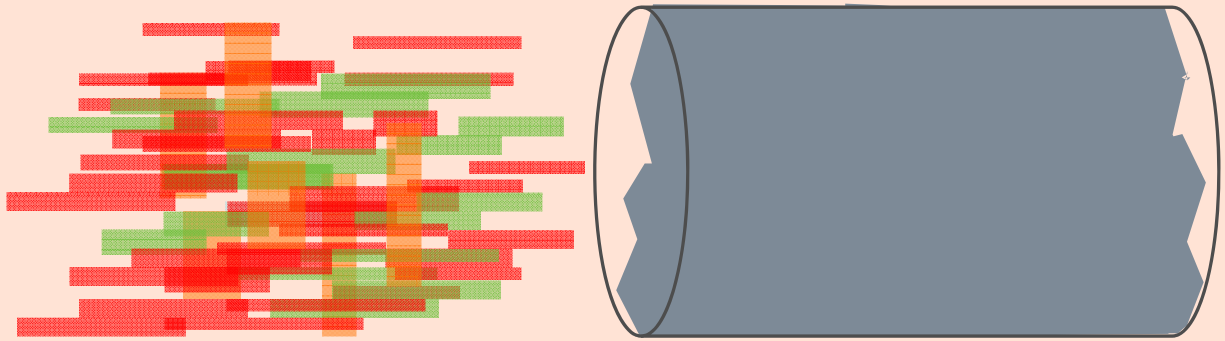

- offered load exceeds capacity ($\rho > 1$)
  - elastic flow throughput $\rightarrow 0$
  - and/or high streaming flow loss
- need for overload control
  - discriminate against some flows
  - discriminate against some classes of traffic

Performance in overload

- processor sharing models are unstable ($\rho > 1$)
  - number of flows in progress increases indefinitely
  - flow throughput tends to zero
- in practice, quality degradation is mitigated
  - by adaptable applications, user impatience
  - by the slow onset of congestion for heavy tailed flow sizes
  - by the presence of non-elastic flows that suffer loss
- but some safeguard seems essential
  - admission control or selective flow termination

Outline

- Traffic and congestion
- Myths and mystification
- Managed or unmanaged Internet

Class of service differentiation

- myth
  - methods exist to realize SLAs of designated classes
- mystification
  - three colour token bucket traffic classifiers…
  - complex class-based schedulers in routers…

Service Class	Service Characteristics	CoS ID	Bandwidth Profile per EVC per CoS ID	Service Performance
Premium	Real-time IP telephony or IP video applications	6, 7	CIR > 0 EIR = 0	Delay < 5ms Jitter < 1ms Loss < 0.001%
Silver	Bursty mission critical data applications requiring low loss and delay (e.g., Storage)	4, 5	CIR > 0 EIR ≤ UNI Speed	Delay < 5ms Jitter = N/S Loss < 0.01%
Bronze	Bursty data applications requiring bandwidth assurances	3, 4	CIR > 0 EIR ≤ UNI Speed	Delay < 15ms Jitter = N/S Loss < 0.1%
Standard	Best effort service	0, 1, 2	CIR=0 EIR=UNI speed	Delay < 30ms Jitter = N/S Loss < 0.5%

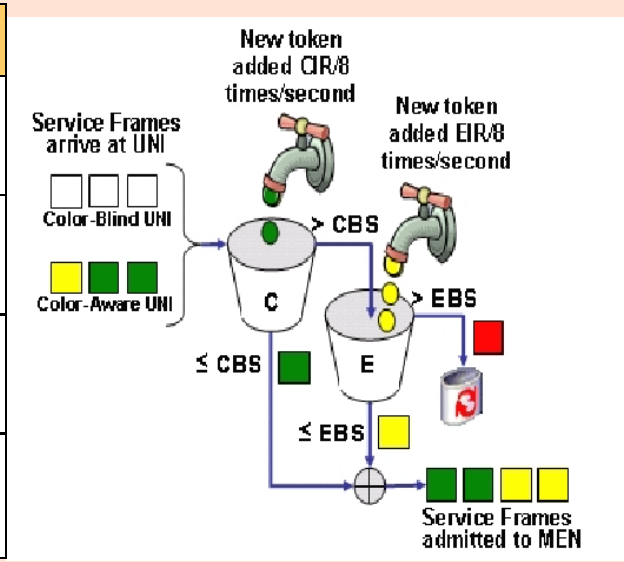

Differentiation is useful to preserve the quality of priority traffic

- priority traffic sees a transparent regime
- implement differentiation by
  - priority queuing, loss priorities, class-based queuing, etc.
- but how is the priority traffic determined?
  - what criterion: application performance or application owner?
  - who decides: user or network operator?

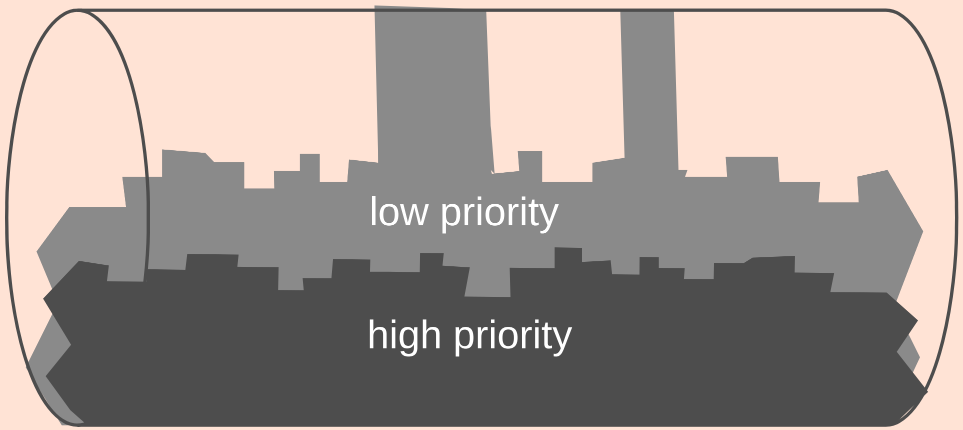

Limited scope for performance guarantees

- can effectively distinguish latency critical and throughput critical applications
  - e.g., conversational & streaming vs interactive & background data
- but no scope to distinguish degrees of guaranteed latency or throughput
  - excellent quality at normal load, too bad in overload
- who can drive the complex class-based router schedulers?
  - a source of mystification!
- no way to reconcile conflicting priority criteria
  - eg, conversational services of low resilience users
  - eg, managed VoD download services

An alternative to class-based: resource allocation and per-flow QoS guarantees

- the principle of ATM (and Intserv, and IMS):
  - user declares the flow "profile"
  - network performs admission control and allocates resources
  - network polices user traffic

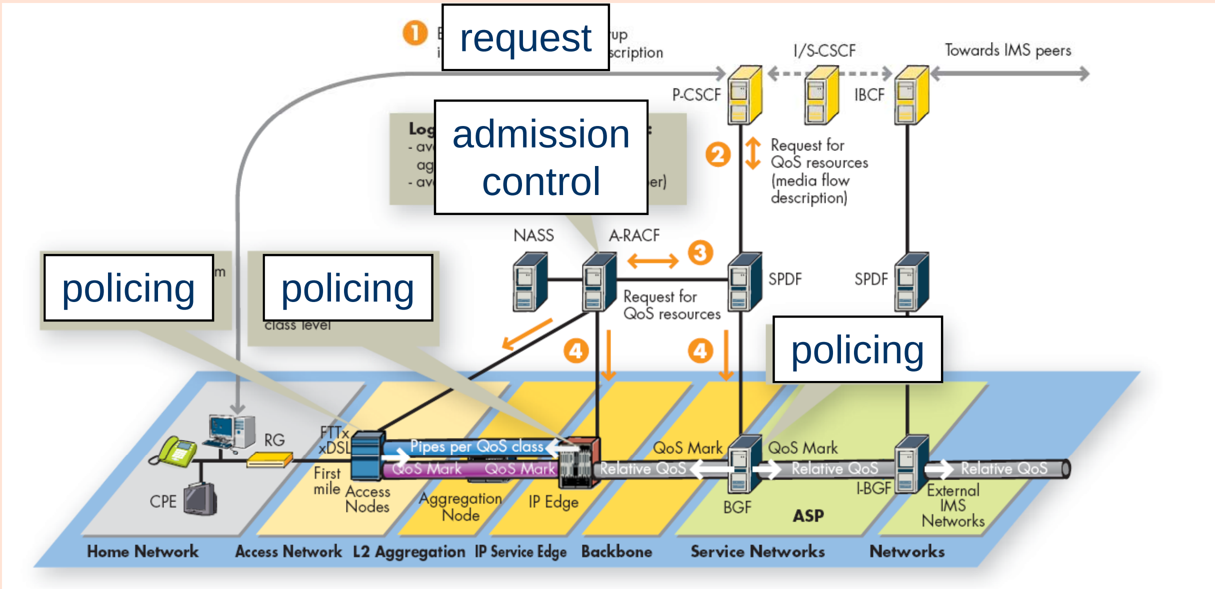

An alternative to class-based: resource allocation and per-flow QoS guarantees

- the principle of ATM (and Intserv, and IMS):
  - user declares the flow "profile"
  - network performs admission control and allocates resources
  - network polices user traffic
- was this ever viable? is it viable for the future Internet?
  - for individual flows or traffic aggregates (eg, for virtual networks)
  - how do we describe the traffic profile?
  - how do we perform admission control?

Reservations and virtualization

- traffic isolation can be achieved by reserving constant rate pipes and implementing scheduling
- but is this satisfactory?
  - for private networks
  - or for service isolation
- seek rather an illusion of isolation based on intelligent, controlled statistical resource sharing via a new QoS paradigm!

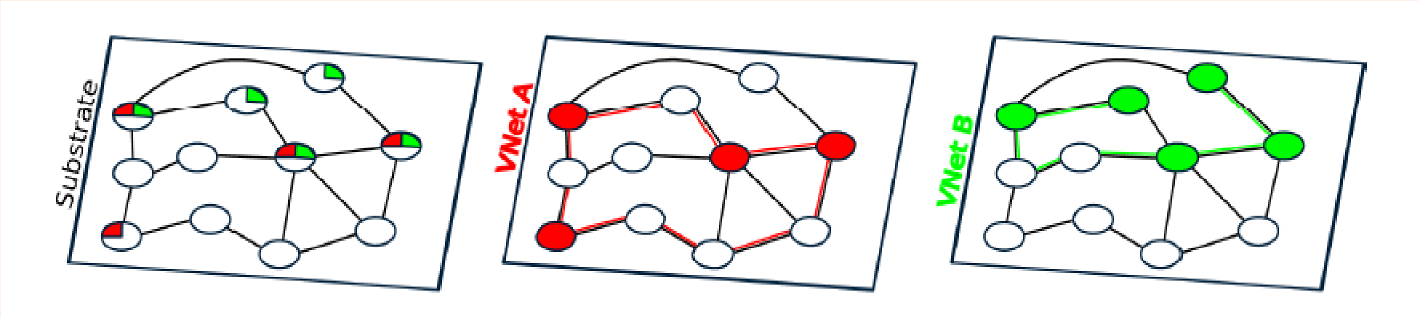

Outline

- Traffic and congestion
- Myths and mystification
- Managed or unmanaged Internet

An operator vision of the future Internet: IPX

- IP interconnection "key benefits"
  - End To End QoS
  - Secure Network
  - Sustainable Commercial Model
  - Universal Interoperability
  - Highly Efficient and Scaleable

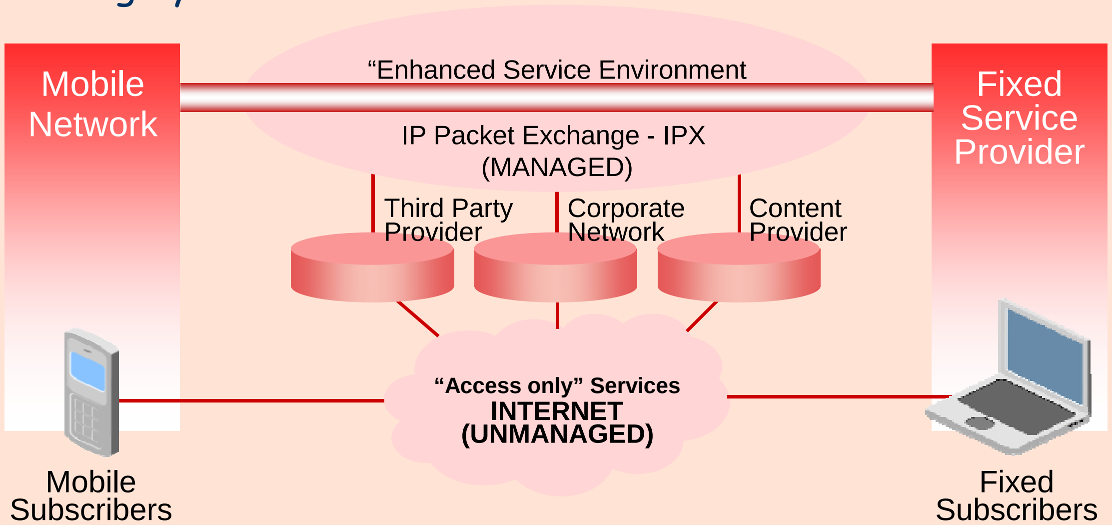

Participants in the IPX trial

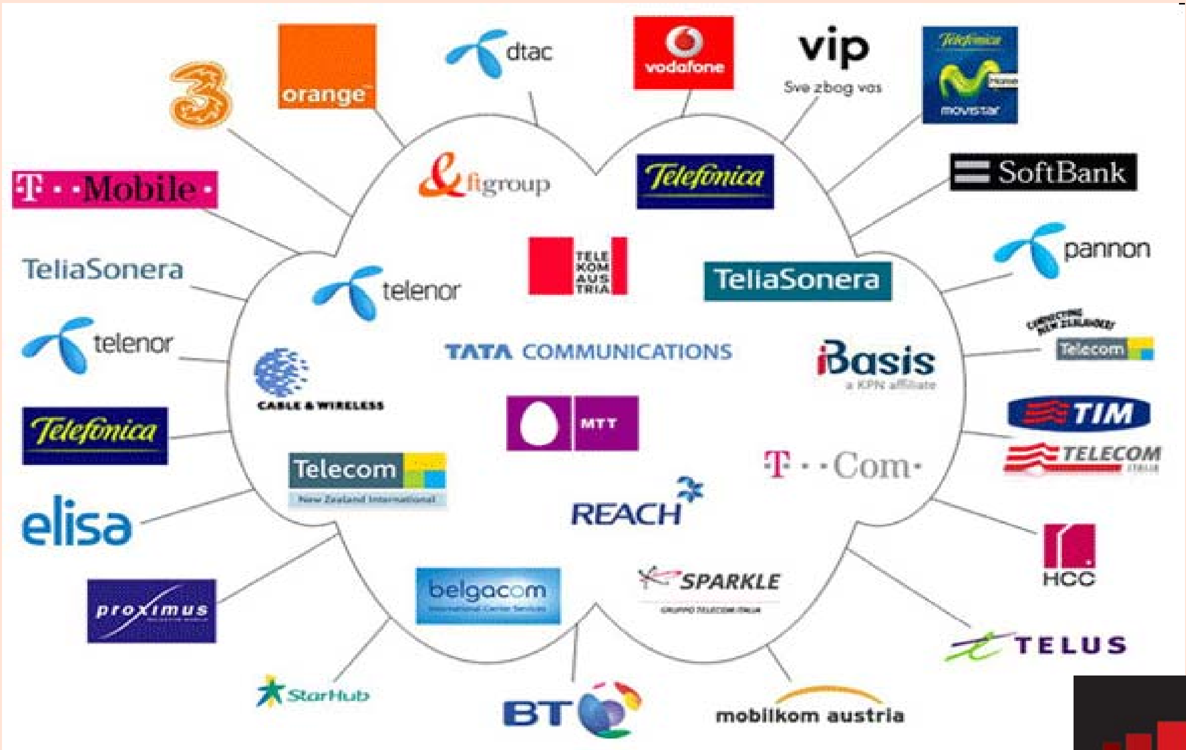

GSM ASSOCIATION

QoS in IPX

- six QoS classes distinguished by DiffServ codepoints
  - conversational, streaming, interactive 1/2/3, background
  - distinct values of delay, jitter, packet loss
- an SLA per class
  - performance is guaranteed and is integral to the business model
- "achieving the QoS specified for the service in the SLA is a matter for the IPX provider"

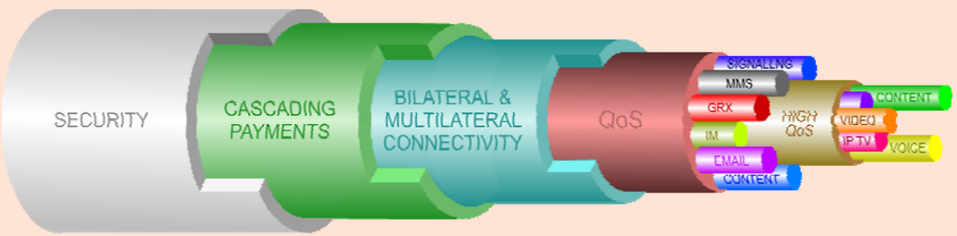

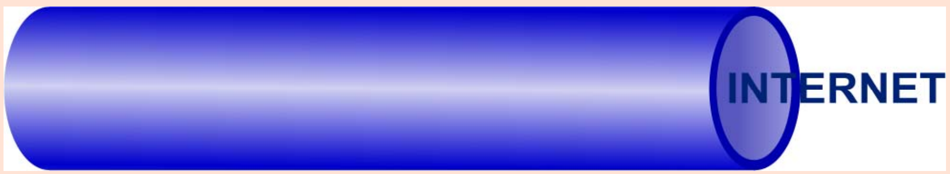

QoS in an unmanaged Internet

- the best efforts Internet works very well most of the time ("thanks to over provisioning")
  - the transparent regime ensures excellent quality for all
- and classical QoS models don't work
  - so much "myth and mystification"
  - driven by the fear of commoditization?
- seek therefore to enhance the best efforts architecture
  - control sharing to prevent abuse and facilitate rate control
  - prevent congestion collapse in overload
- and allow end users to control sharing of last mile resource

Network controlled bandwidth sharing

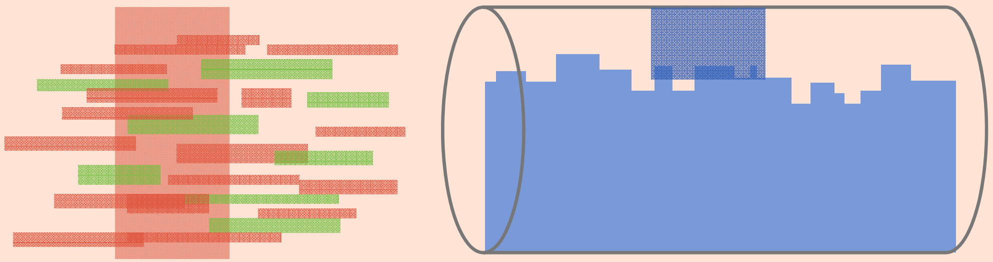

- impose per-flow fair sharing in router queues
  - using fair queuing (or just fair dropping?)
- this realizes implicit differentiation
  - no loss for flows of rate < fair rate, i.e., all streaming flows
- fair queuing is scalable (cf. Kortebi, et al., 2005)
  - only $O(10^2)$ flows (with a packet in queue) need scheduling
- overload control is necessary (exceptionally)
  - flow admission control or selective flow termination to maintain the fair rate high enough
- of course, users can disguise high rate flows! (cf. Briscoe, 2007)
  - but we only need to identify the *very* high rate flows

User control over last mile resources (DSL, wireless, fibre,...)

- in current networks, the operator dictates its own priorities, even upstream in the user's home gateway
  - e.g., priority to managed services that earn revenue!
- but only the user knows its own preferences
  - based on the application but also on the end user, time of day,...
- proposal: user signals priority for each flow, access node implements (same as for upstream traffic in gateway)
  - e.g., priority to Skype, priority to Dad,...

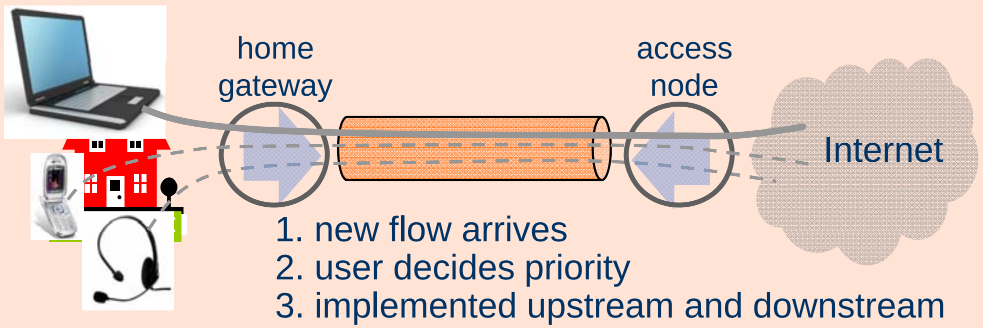

Conclusion: we still need research on QoS for the future Internet '!

- operators would seek to create a service-based Internet with architectures like IPX and IMS
  - for added value services and vertical integration
- but these architectures rely on QoS models built on "myth and mystification" rather than sound engineering
  - neither class-based differentiation nor per-flow guarantees make much sense as traffic controls
- prefer a neutral, flow-aware future Internet
  - network controlled sharing of high capacity links for implicit differentiation and flexible congestion control
  - implement controls to avoid congestion collapse in exceptional overloads
  - user controlled sharing of last mile resources, upstream and downstream
- towards communications networking as a utility?