What QoS for the future Internet?

Jim Roberts
ex France Telecom (soon INRIA)

FISS 2009
University of Bremen
22 July 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
  - enhanced network management
- 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!

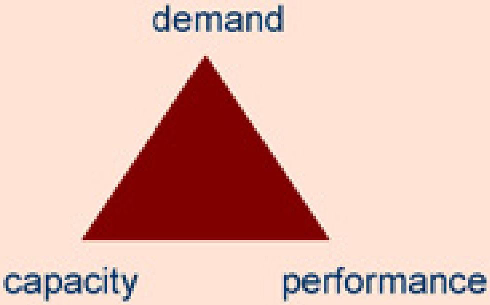
demand

capacity

performance

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

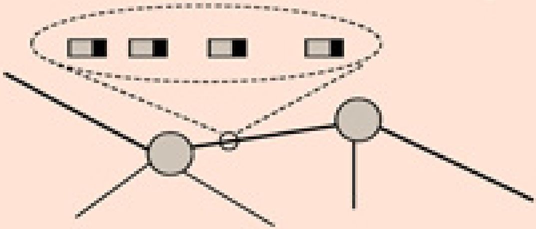

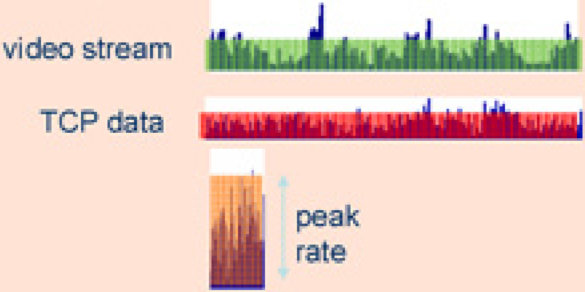

Two essential traffic characteristics

1. the mix of flow rates

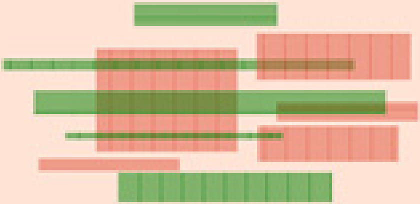

2. the link load

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

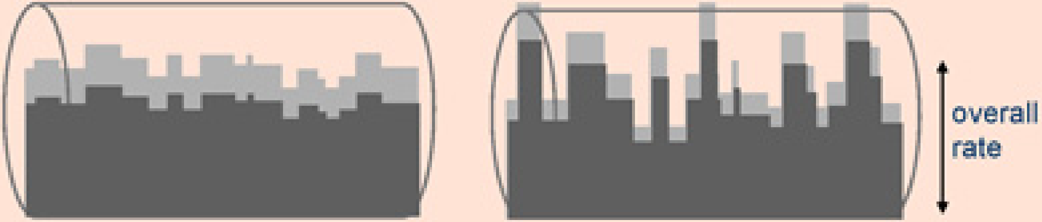

Transparent bandwidth sharing

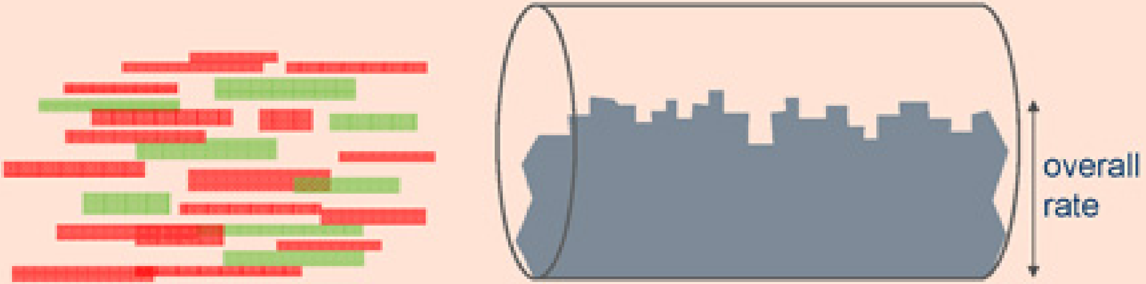

- all flows have relatively small rate
- offered load is somewhat less than capacity ⇒ no rate excess
- excellent quality for all
- this is "QoS by over provisioning"

Performance in the transparent regime

- predictable packet level performance:
  - flows combine to produce a "better than Poisson" arrival process
  - $\Rightarrow$ a modulated *M/G/1* queue is a conservative model
- negligible delay at normal load (<90%)
  - e.g., one 1500B packet at 1 Gb/s $\Rightarrow$ 12μs
- performance is insensitive to detailed traffic characteristics
  - e.g., to self-similarity in flow or aggregate arrival processes
  - not true if overall rate can exceed capacity!

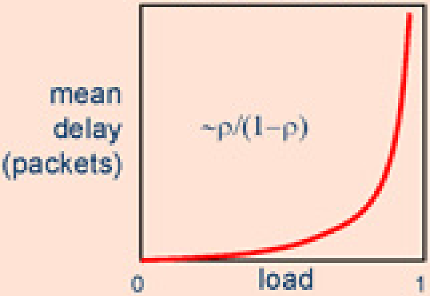

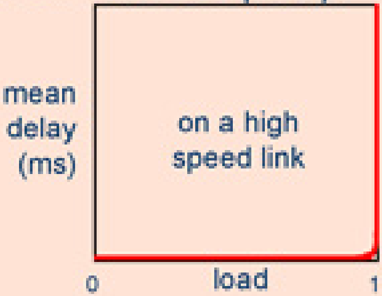

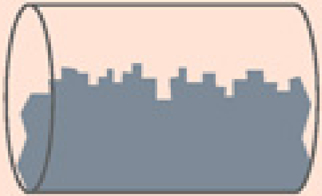

Maintaining the transparent regime

- max utilization depends on link rate to peak rate ratio
- this is typically high for most flows on shared links
- but some high rate flows may occur: inter-server flows, physics labs,...

max utilization

link rate / peak rate	Pr [excess rate]	
	0.01	0.001
10	48%	38%
100	79%	73%
1000	92%	90%

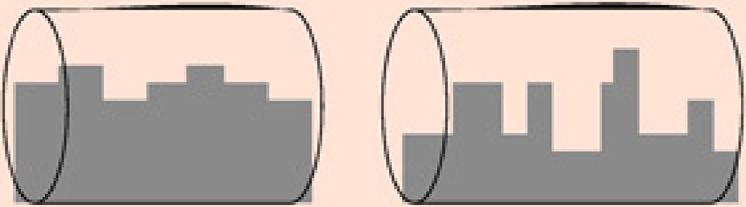

Elastic bandwidth sharing

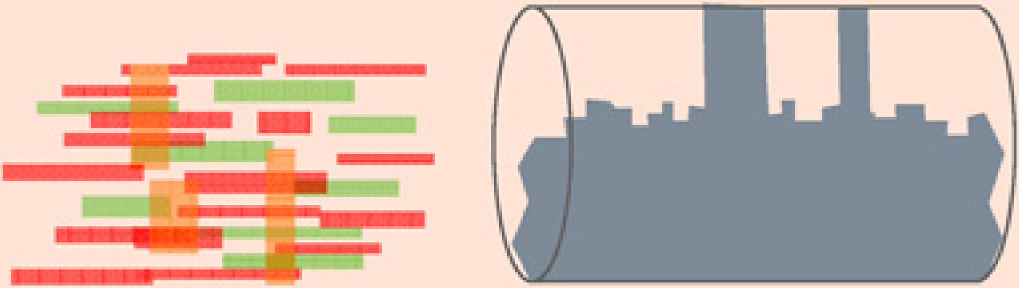

- some flows have high rate
- offered load somewhat less than capacity
- need to control sharing
  - to avoid loss to low rate flows
  - to ensure "fair" sharing

Performance in the elastic regime

- a "processor sharing" model of bandwidth sharing
  - assume all flows of unlimited rate, Poisson flow arrivals, instantaneous fair sharing
  - $\Rightarrow$ *M/G/1* processor sharing model
    - Pr [n flows] = $\rho^n (1 - \rho)$
    - E [response time] = size/$C(1 - \rho)$
- same results apply for more general traffic
  - Poisson <u>session</u> arrivals, alternating sequence of flows and think times (with general distributions, correlation)

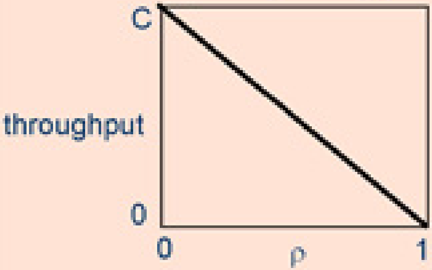

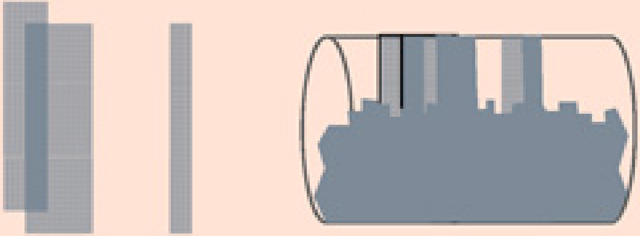

Performance in the elastic regime (2)

- extension to unfair sharing (discriminatory PS) cf. Fayolle et al. 1980
  - (somewhat) sensitive performance
  - little discrimination at normal loads
- accounting for peak rate limit
  - performance for equal rate or "balanced fair sharing"
  - E [throughput] ≈ min {peak rate, $C(1 - \rho)$}
- performance typically excellent, depending only on ρ
  - for more info, cf. Bonald et al., Queuing Systems, 2006

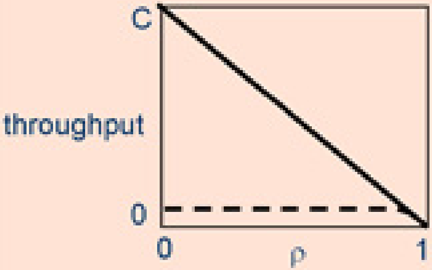

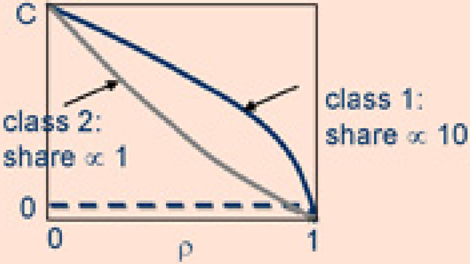

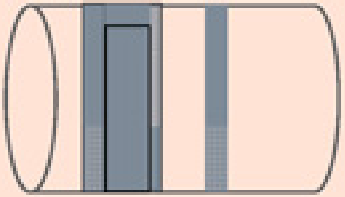

Controlling the elastic regime

- if elastic flows are suitably responsive to congestion, quality of low rate streaming flows may be adequate
  - small buffers and ECN to limit delay
  - but what incentive for users to be "TCP friendly"?
- economic incentives for congestion control
  - eg, a "self-managed Internet" (F. Kelly, 2000):
    - a charge for each ECN mark, reduced rate ⇒ reduced charge
- but congestion pricing is unacceptable, prefer network imposed fairness?

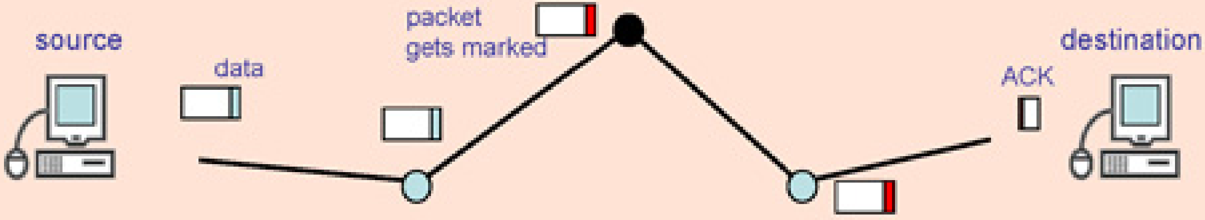

Overload and bandwidth sharing

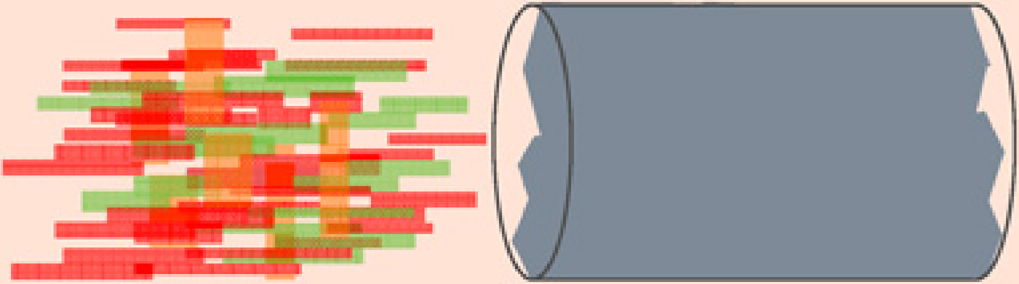

- offered load exceeds capacity ($\rho > 1$)
  - elastic flow throughput → 0
  - and/or high streaming flow loss
- need for overload control
  - discriminate against some flows
  - discriminate against some classes of traffic

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

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...

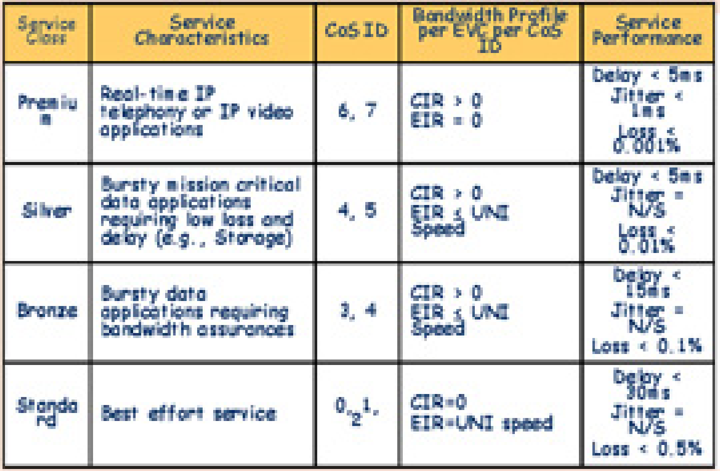

Service Class	Service Characteristics	CoS ID	Bandwidth Profile per EVC per CoS ID	Service Performance
Premium	Real-time IP telephony or IP video applications	6, 7	CIR > 0 EIR = 0	Delay < 5ms Jitter < 1ms Loss < 0.001%
Silver	Bursty mission critical data applications requiring low loss and delay (e.g., 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%

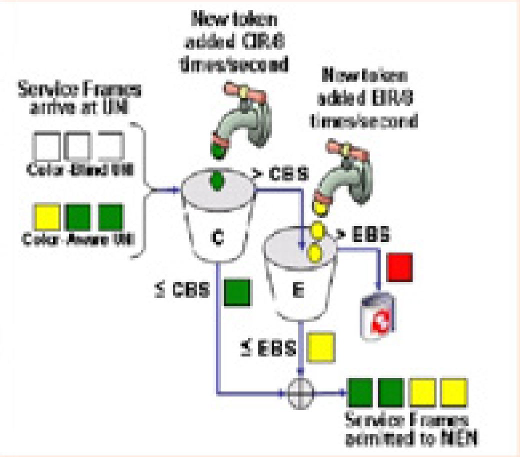

Differentiation is useful to preserve the quality of priority traffic

- 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, who decides

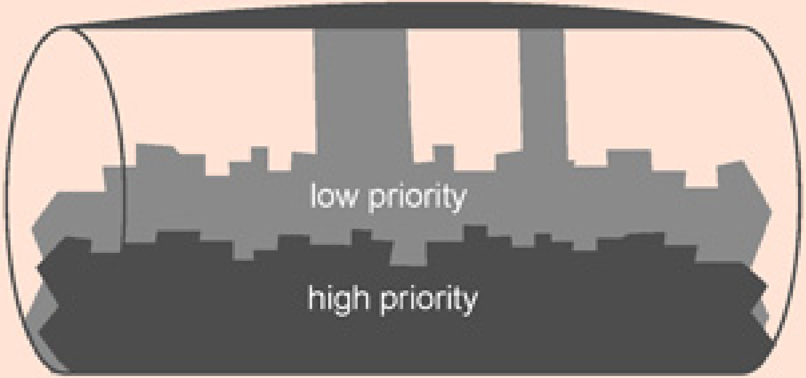

Multiple criteria for defining priorities

- differentiating application requirements
  - conversational, streaming, interactive data, background data
- differentiating services and/or users
  - priority to operator's "managed services"
  - tiered service offering to other providers
- differentiating user reliability requirements
  - "five 9s" reliability for business subscribers
- differentiating applications
  - filtering P2P traffic
  - no VoIP on 3G wireless

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

Performance depends on demand and capacity

- e.g., the M/M/1 queue
  - E [delay] = $\tau \rho / (1 - \rho)$, τ = packet time, ρ = link load

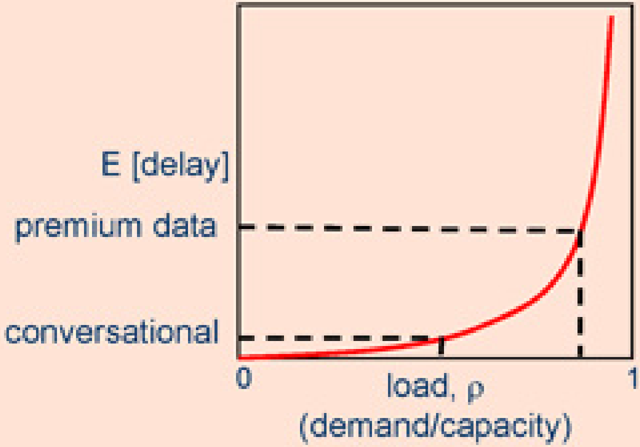

Delay/Load Tradeoff

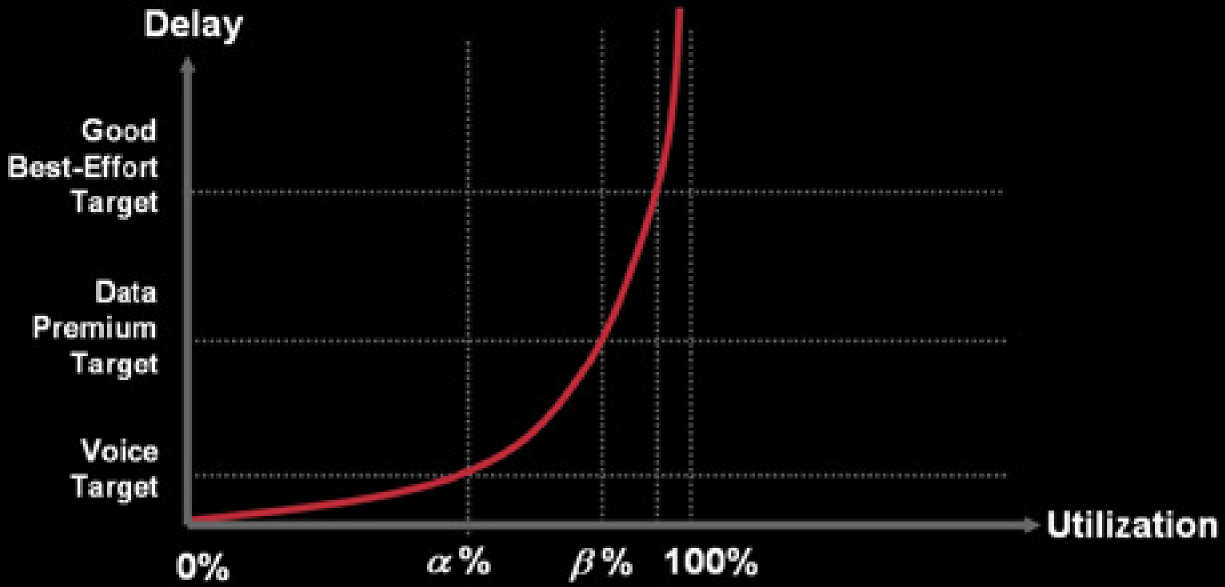

If I Can Keep EF Traffic < α%, I Will Keep EF Delay Under *M1* ms
If I Can Keep AF1 Traffic < β%, I Will Keep AF1 Delay Under *M2* ms

Performance depends on demand and capacity

- e.g., an *M/M/1* queue
  - E [delay] = $\tau \rho / (1 - \rho)$, τ = packet time, ρ = link load
- very little scope for service differentiation
  - quality of service is "good" or "bad"
- a need for overload control (when $\rho \geq 1$)

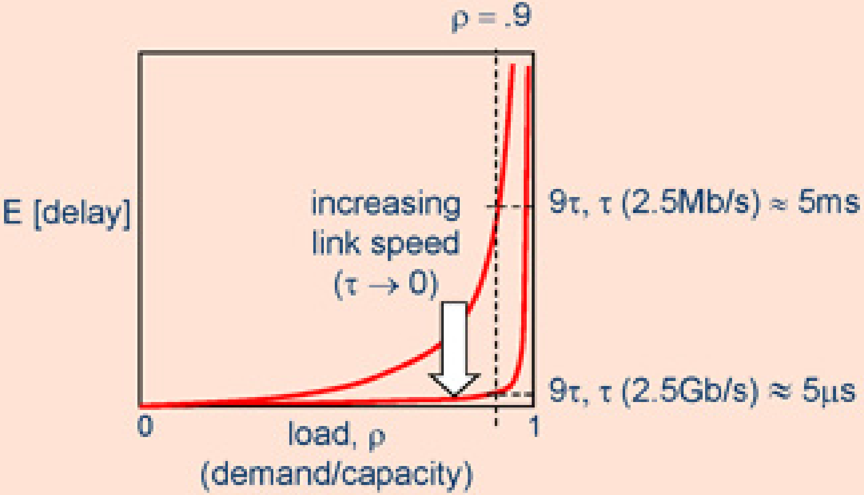

Limited scope for performance guarantees

- 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
- need to account for conflicting priorities
  - eg, conversational services of low resilience users
  - eg, managed VoD download services
- difficult to reconcile using class-based differentiation
  - we would need to satisfy SLAs for a matrix of classifications
  - routers implement complex class-based schedulers that are in practice uncontrollable (mystification!)

Class-based differentiation and network neutrality

- differentiation is unfair and stifles innovation
  - eg, priority to managed services, favouring vertical integration
  - eg, tiered services ⇒ quality for the rich and powerful
  - eg, deep packet inspection to discriminate against applications
- should we expect regulatory constraints?
  - network neutrality and the US congress
  - future European directives on neutrality
  - separation of infrastructure and service provision
- NB. neutrality might be an advantage for the operator
  - simpler operation in the absence of classes
  - meet user requirement for choice (e.g., priority to Skype)

An alternative to class-based: per flow QoS guarantees

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

QoS in IMS

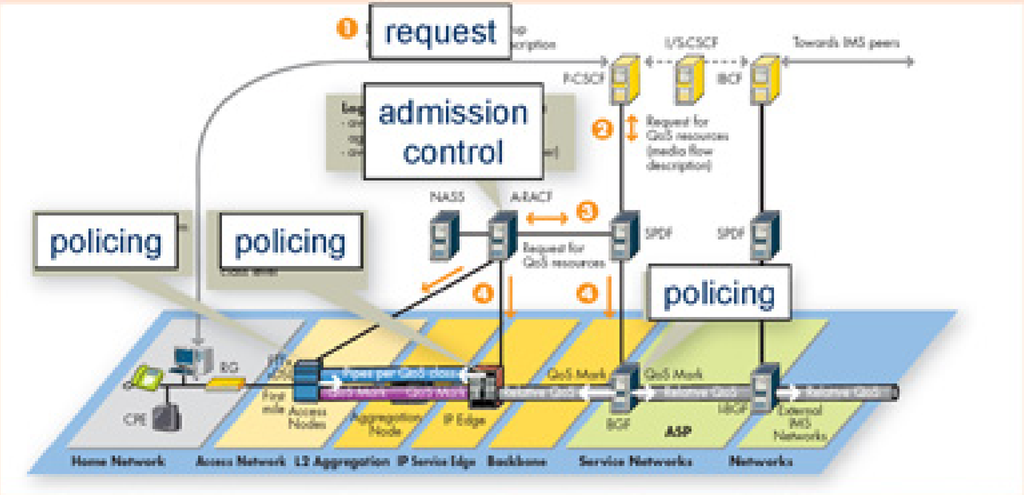

Figure from Alcatel-Lucent white paper: Quality of Service for IMS on Fixed Networks

- but, how do we describe the traffic profile?
- how do we perform admission control?

Admission control for variable rate flows

- traffic descriptors should be
  - understandable, meaningful, controllable
  - cf. ITU Rec I.371 (~1990)
- for practical reasons, only the latter requirement is satisfied
  - by the leaky/token bucket!
- admission control is then either
  - too conservative, based on worst case traffic assumptions
  - or imprecise, based on empirical over booking factors
- the only satisfactory solutions are measurement based algorithms designed to maintain the transparent regime

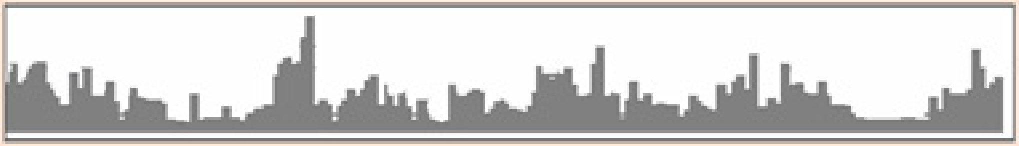

Reservations and virtualization

- 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!

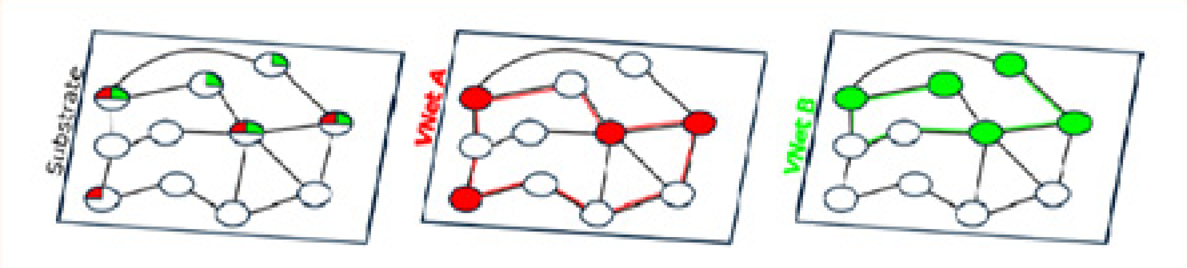

Outline

- 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

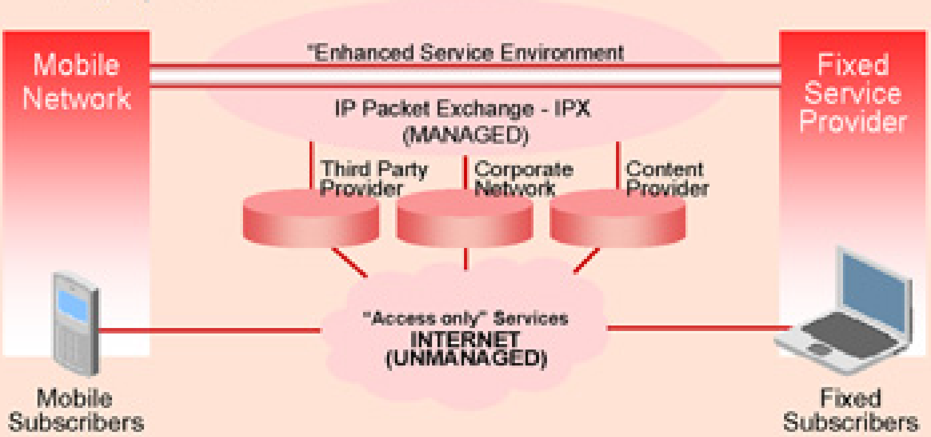

Participants in the IPX trial

GSM ASSOCIATION

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

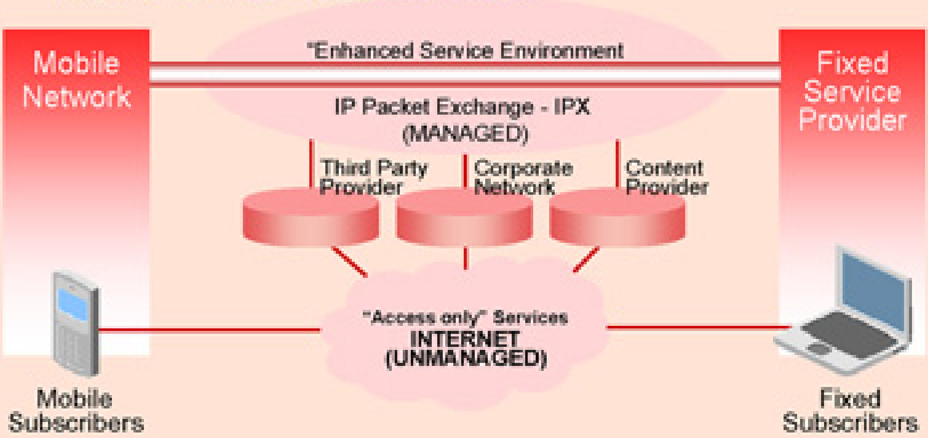

QoS in IPX

- 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"

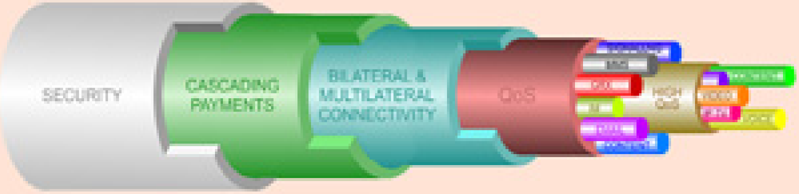

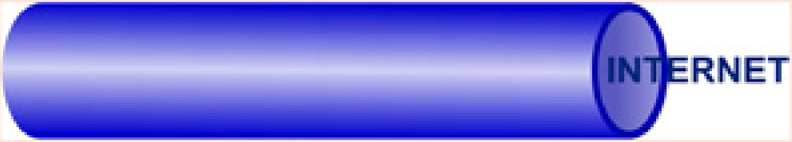

QoS in an unmanaged Internet

- 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 flow control
  - prevent congestion collapse in overload
- and allow end users to control sharing of last mile resource

Network controlled bandwidth sharing

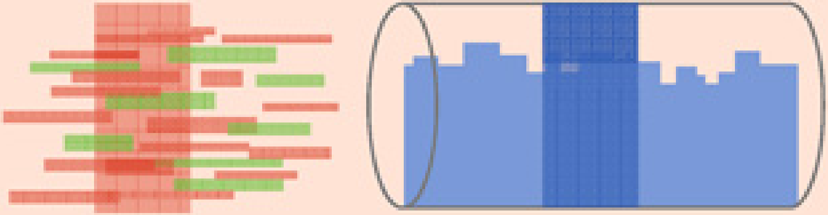

- impose per-flow fair sharing in router queues
  - using fair queuing (or just fair dropping?)
- this realizes implicit differentiation
  - no loss for flows of rate < fair rate, i.e., all streaming flows
- fair queuing is scalable (cf. Kortebi, et al., 2005)
  - only $O(10^2)$ flows (with a packet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

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,...

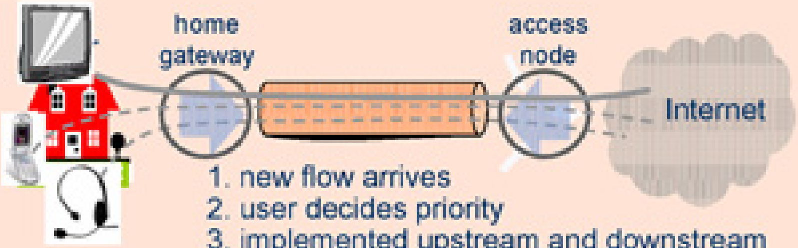

1. new flow arrives
2. user decides priority
3. implemented upstream and downstream