Capacity Sharing

Bringing Together Cost, Value and Control

Bob Briscoe
Chief Researcher
BT
Oct 2011

This work was partly funded by Trilogy, a research project supported by the European Community
www.trilogy-project.org

a talk in two parts

1. sharing capacity between commodity traffic
   - the network side (ensure end-to-end transports are nice)
   - net neutral for ‘over-the-top’ services
   - looks like engineering, but based on economics

2. sharing capacity with added value services
   - change gear – commercial, but based on engineering

- mobile networks were designed for #2
  - treating #1 as an afterthought was a mistake

Capacity Sharing: Lest we Forget

- sharing is central to all developments in network access
  - and obviously core, campus & enterprise networks too
- for dedicated consumer access, utilisation ⇘ as speed ⇗
  - ave. utilisation during peak 15min only 0.5% for 40M access
  - was 1.25% for 4M access; still higher in dial-up days
- cost efficiency will drive sharing closer to the end-user
  - that's why we see cellular, cable, passive optical networks, WiFi
  - that's why we see packet multiplexing & virtualisation
    - IP, Ethernet, MPLS
- central dilemma
  - performance isolation without losing efficiency of multiplexing

Internet topology visualization produced by Walrus (Courtesy of Young Hyun, CAIDA)

existing approaches are hopeless

‘fair’ bandwidth

- (weighted) equal bandwidth per-active-user
  - (weighted) round robin per-user
  - (weighted) fair queuing per-user
- each of *n* active users gets $^1/_n$ of capacity

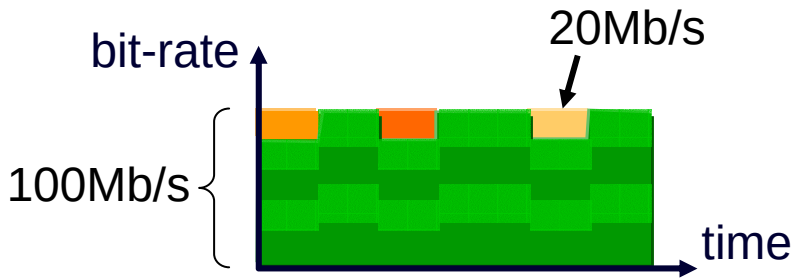

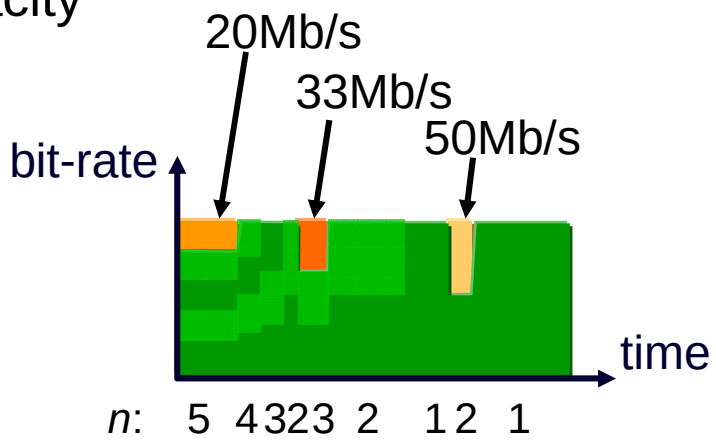

- how much I get
  - highly dependent on how often others are active
- no concept of time; no memory

existing approaches are hopeless

accounting by volume

- introduce memory – a per-customer account
  - e.g. RADIUS or a traffic management node
  - measurable at one control point
  - not just per-link like WRR, WFQ
- but …

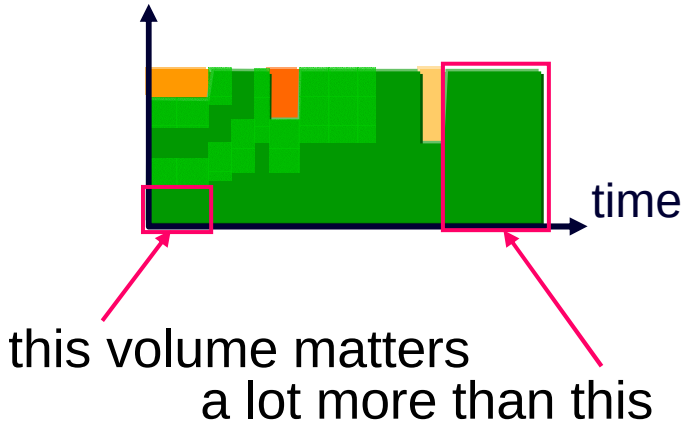

- simple sum of volume loses information about
  - conditions at each link (space)
  - how conditions change (time)

existing approaches are weak

peak period volume

- state of the art in traffic management
  - count only peak-period volume
    - Comcast Fairshare [RFC6057]
    - some deep packet inspection (DPI)
- better, but...
- ...still penalises ‘ultra-friendly’ volume
  - BitTorrent [µTP] & LEDBAT
  - equitable quality streaming (next slide)
- p2p & video
  - could potentially deliver a large proportion of traffic
  - while minimising impact on interactive apps

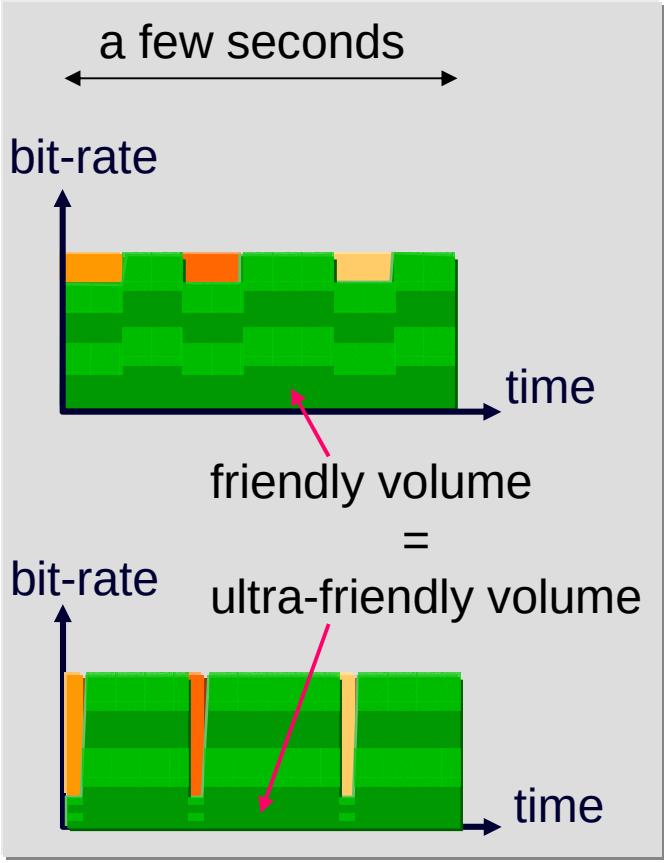

equitable quality streaming (EQS) video [Mulroy09]

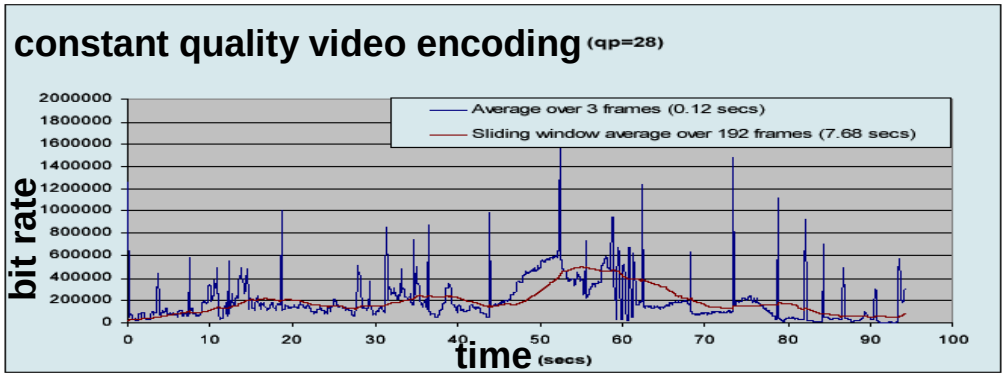

- near-constant mean opinion score
  >2x more videos in same capacity

- delivered over MulTCP
  - TCP with weight parameter, n
  - adjust n to 'hardness' of video

  - if peaks coincide
    - all MulTCPs back-off
    - whether in a peak or trough
    - equitable loss of quality
    - even if one taking more bit-rate than another

- in contrast 'fair' queuing forces all EQS streams to have equal bit-rate
  >2x less videos in same capacity ☹

summary so far

- multiple traffic management approaches deployed
  - fair queuing
  - volume limits
  - peak-volume-based traffic management
- often all three, and more… (QoS mechanisms)
- each patched a problem of its time
- a symptom of never really understanding the problem

measuring
contribution to congestion

- user's contribution to congestion
  = bytes weighted by congestion level
  = bytes dropped (or marked)*
  = 'congestion-volume'
- as simple to measure as volume

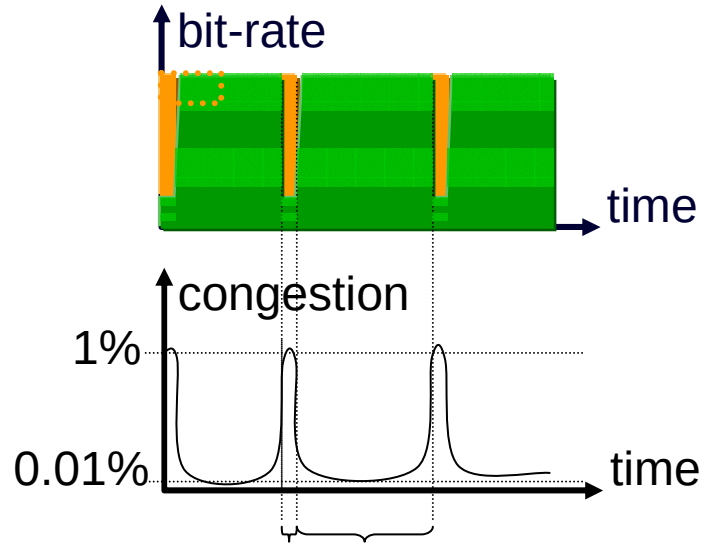

---

* congestion [Siris02] & Kutsher (next)
- radio uplink (interference)
- fixed links (queue lengths)
- radio downlink (power)

all signalled up and along the IP layer

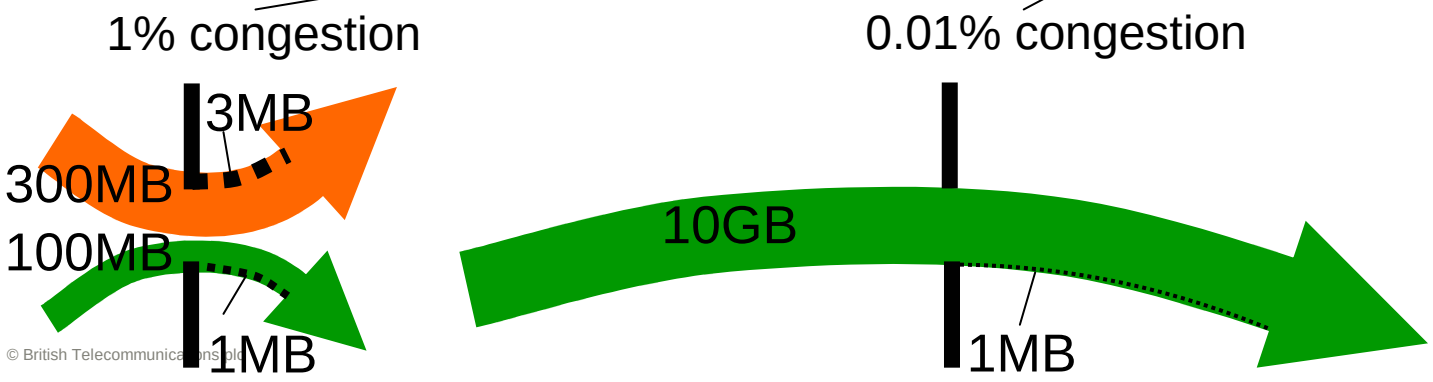

if only... ingress net could see congestion...

flat fee congestion policing

Acceptable Use Policy

'congestion-volume' allowance: 1GB/month

@ €15/month

Allows ~**70GB per day** of data in typical conditions

- performance isolation without losing efficiency of multiplexing
- edge control point
- controls capacity sharing of any link without modifying any links

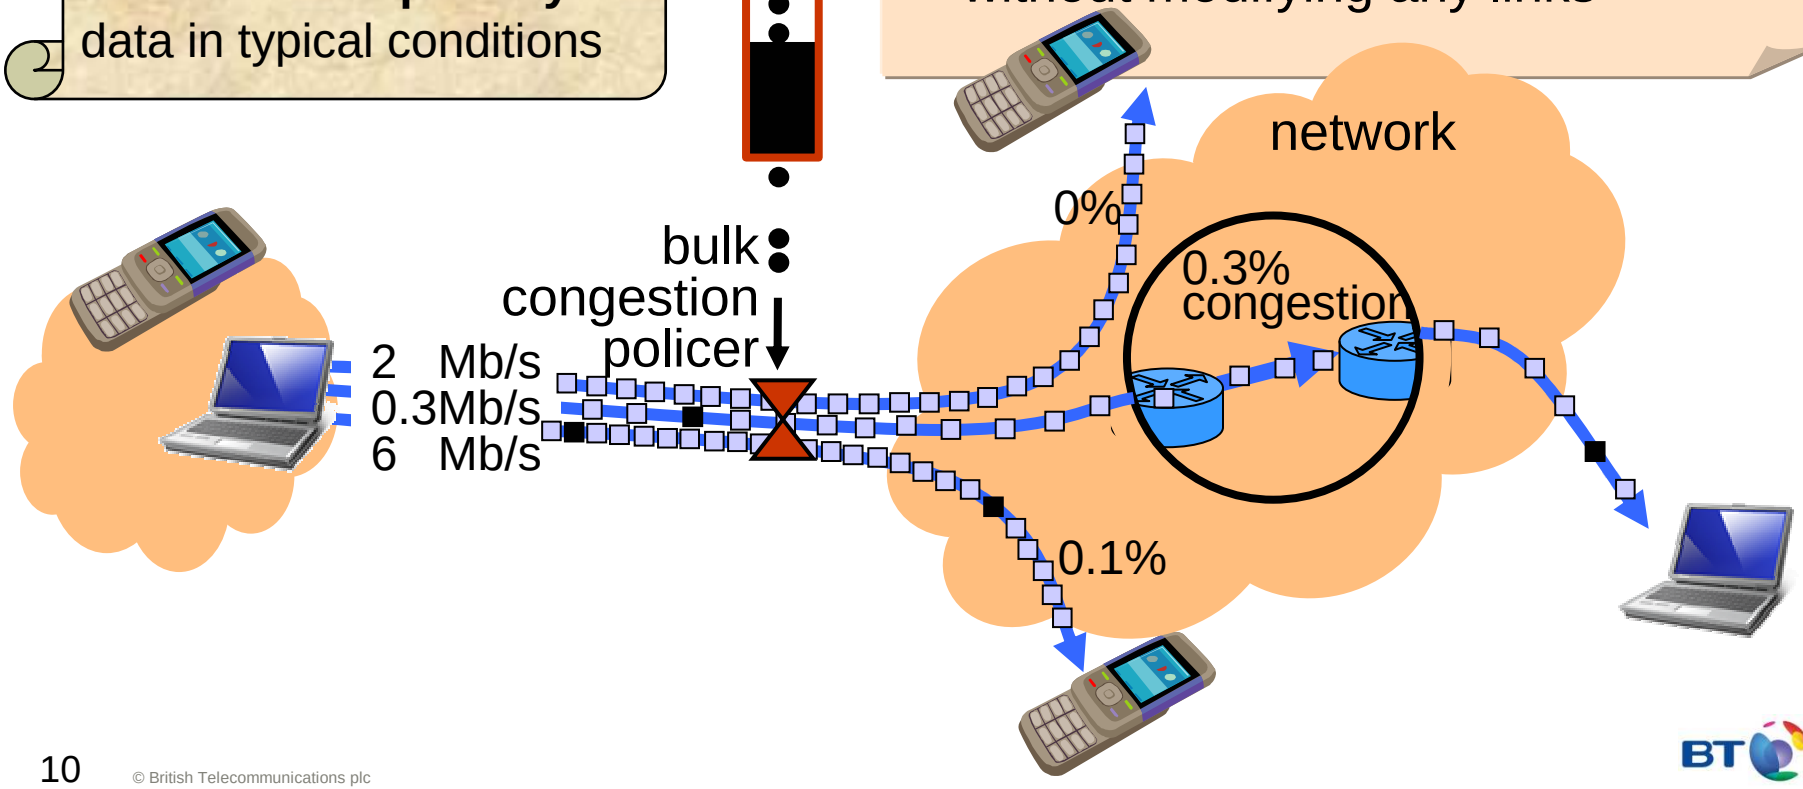

“the capacity sharing I do isn’t about congestion”

- the focus on congestion can be misleading
  1. congestion signals to *avoid* quality impairment (spare slide)
  2. not ‘solving’ or minimising congestion *per se,*
     using congestion signals to control capacity sharing

- for instance
  - if many users send continuously through one link
  - outcome of equal congestion policing would be equal bit-rates

- exploiting an inescapable fact
  - the greater the share of capacity you use
    when others would like to use more,
    the more congestion signals will be attributed to you

---

nonetheless, with ECN and well-sized thresholds in buffers and radio resources

- congestion policing would maintain low latency and low loss for all too

“if only... ingress net could see congestion…”

Congestion Exposure [ConEx]

IETF activity to add new capability to IP

drop or standard ECN + re-inserted feedback (re-feedback)

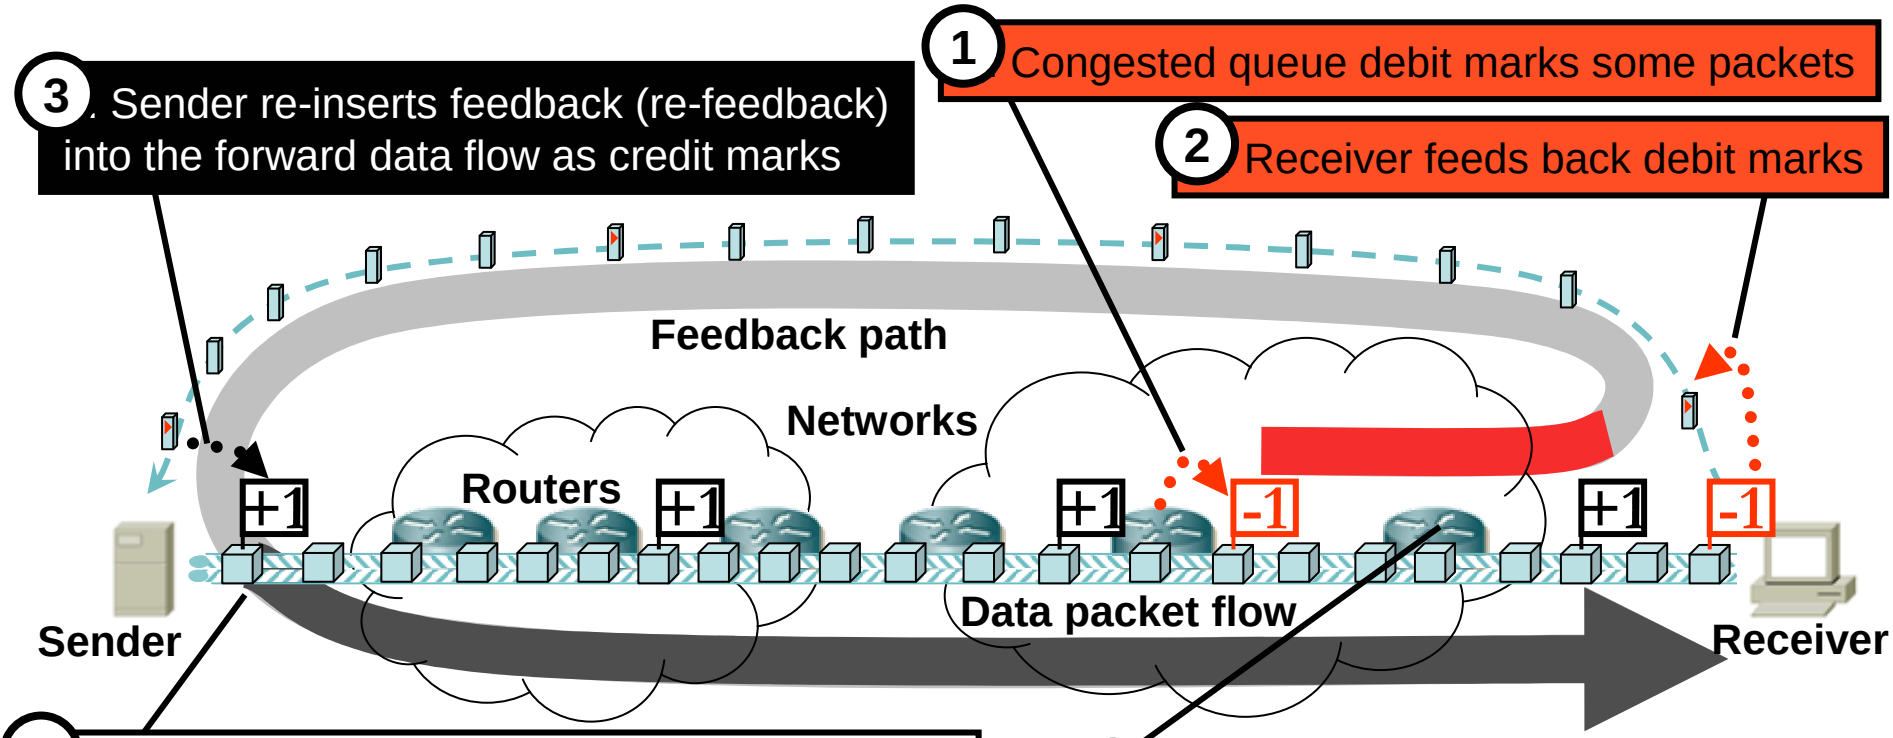

no changes required to IP data forwarding

both value *and* cost

- even a CEO should understand both value *and* cost
- maximise profit = revenue – cost
  1. competitive market drives revenues down towards provider's marginal cost
  2. until then, revenue depends on consumer value

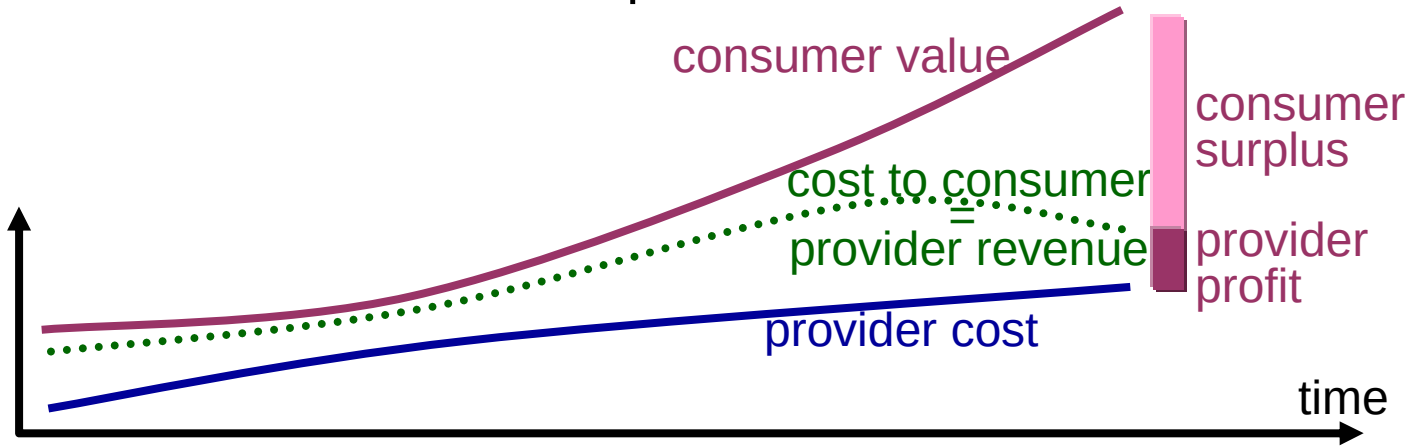

- in both cases
  - those who understand marginal costs will succeed
  - the marginal cost of traffic *is* its congestion-volume

added value and over-the-top services

- today's policy control box (DPI)
- but with real-time cost info
  - brought to it in the packets
- cost info today is inaccessible
  - dynamic and
  - mainly in radio access network

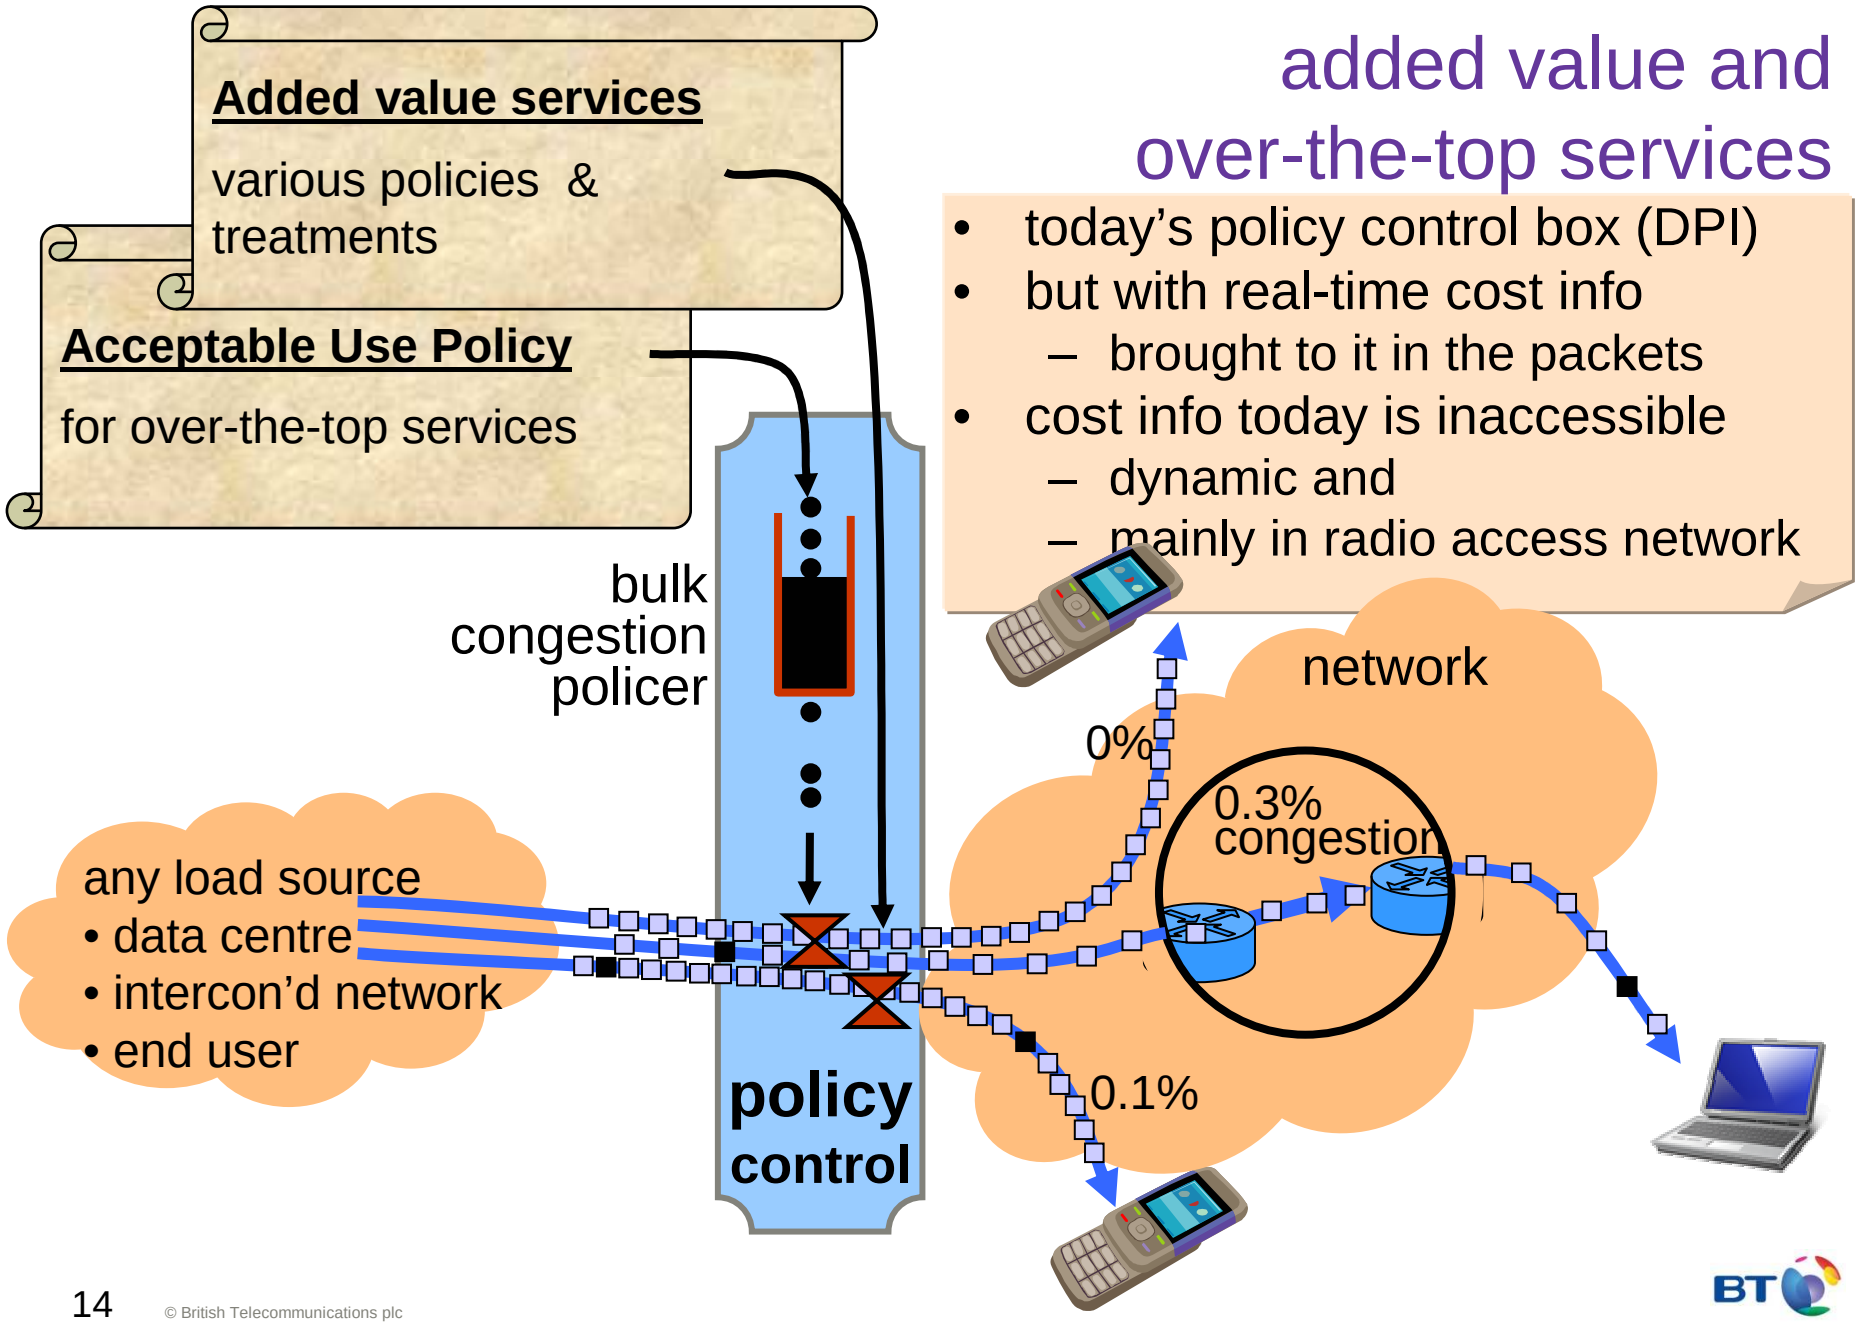

remote viewing from the traffic mgmt box

- the volume of a flow is the same wherever metered
- but, whole path congestion-volume can be split
  - into congestion downstream and upstream of a point
  - can measure all three with ConEx
  - want to ignore congestion on the customer's side
- at policy control box
  - can 'remote-view' split betw upstream & downstream
  - from point where customer attaches to your network
- solution: tunnel from the attachment point
  - inner headers 'freeze' congestion info from tunnel ingress

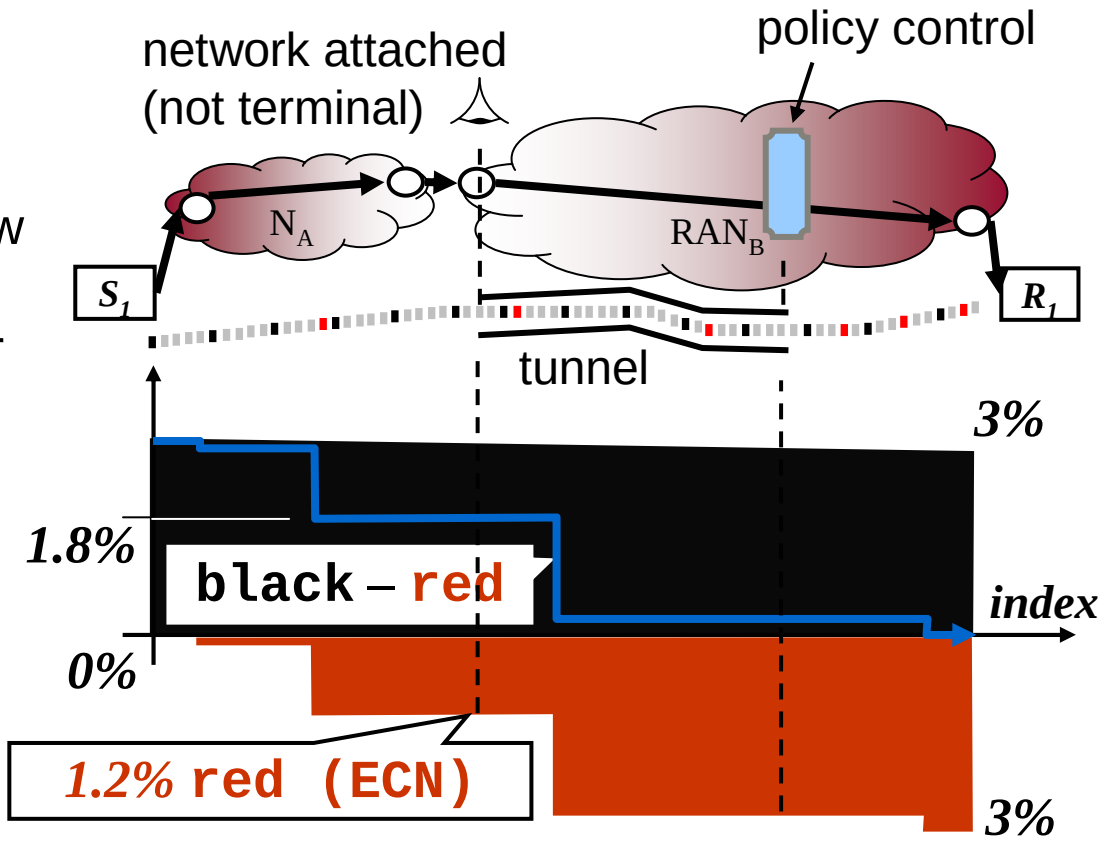

openness can be a tactic not a strategy

- edge policer is the focus all policy enforcement
  - **open:** per-user policing of bulk congestion volume
    - will allow much more freedom to innovate than current FQ constraint
    - new behaviours: e.g. very hi-speed, unresponsive, weighted, networked games
    - but all within overall responsibility envelope
  - **closed:** policing / enhancement of specific applications
    - optimising perceived value against marginal cost
    - using cost information carried in ConEx packets
- MVNOs / Retailers choose
  - how tightly to control true network costs
  - each product's market pos
    between open and closed
- Changing your mind
  - involves changing a policy
  - not new technology
- MNO / Wholesaler is agnosti
  - supports all positions
  - simultaneously

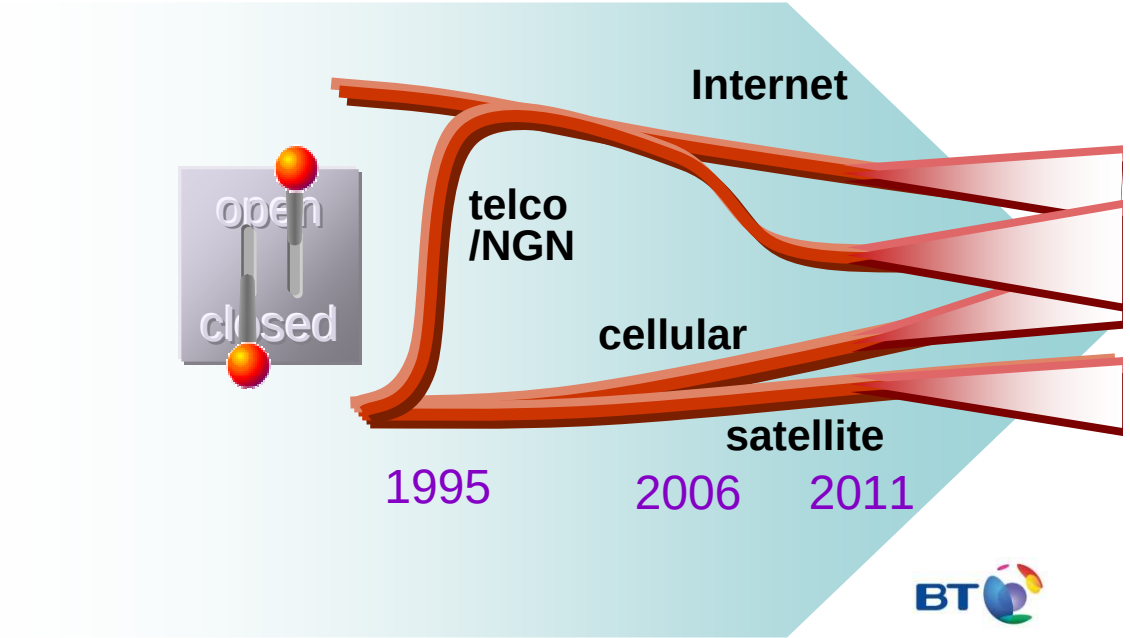

summary

- bringing together cost, value and control
  1. real-time marginal network cost info (ConEx-IP)
  2. market knowledge on customer value (DPI)
  3. ‘edge’ control point (on path, near edge)

- cost info is actually more important than value
  – to handle over-the-top traffic today
  – and as the market commoditises

- information on marginal cost is then all we need
  – but it’s all we haven’t got
  – we’re working on this in ConEx at the IETF

references

[Mulroy09] Mulroy, P., Appleby, S., Nilsson, M. & Crabtree, B., "The Use of MulTCP for the Delivery of Equitable Quality Video," In: *Proc. Int'l Packet Video Wkshp (PV'09)* IEEE (May 2009)

[µTP] Norberg, A., "uTorrent transport protocol," *BitTorrent.org* BitTorrent Enhancement Proposals (BEPs) 0029 (January 2010) (Draft)

[Siris02] Siris, V.A., "Resource Control for Elastic Traffic in CDMA Networks," In: *Proc. ACM International Conference on Mobile Computing and Networks (MobiCom'02)* ACM (September 2002) <http://www.ics.forth.gr/netlab/future_wireless.html>

Congestion Exposure (ConEx) and re-feedback:

<http://bobbriscoe.net/projects/refb/>

Bob.Briscoe@bt.com

BT

Capacity Sharing

Bringing Together
Cost, Value and Control

Q&A

congestion is not evil
congestion signals are healthy

- no congestion across whole path ⇒ feeble transport protocol
  - to complete ASAP, transfers should sense path bottleneck & fill it

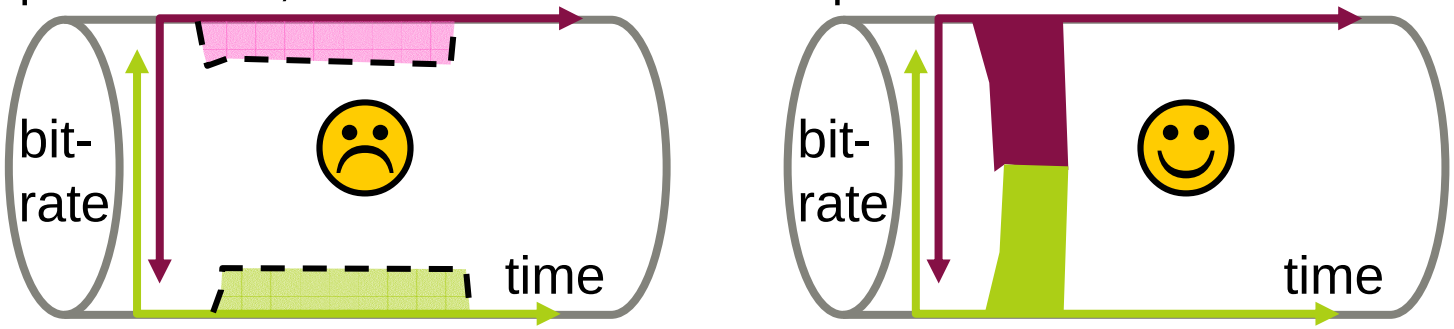

the trick

congestion signal *without* impairment

- explicit congestion notification (ECN); update to IP (2001)
  - mark more packets as queue builds
  - then tiny queuing delay and tiny loss for all traffic
  - no need to avoid congestion signals to prevent impairment
- original ECN: gain too small to overcome deployment barriers