



the speed of sharing stretching Internet access

Bob Briscoe
Chief Researcher, BT
Apr 2009



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



shared access






- India: 11,000 new mobile contracts /hr
- given best available access technology
 - huge gains from sorting out sharing properly
 - currently a disaster area




harness mutual flexibility

- faster when you really need it
- greater value, better quality of experience
- gentler entry ramp to the Internet
 - share infrastructure cost between more people
- inability to prevent free-riding kills capacity investment [CFP06]



how to share a bandwidth cloud?




source: Ellacoya 2007
(now Arbor Networks)

- since 1988: misplaced belief that 'TCP-friendly' sharing is good
 - but ISP's homespun alternatives have silently overridden TCP
- since 2006 IETF support for TCP-friendly sharing has collapsed
 - Van Jacobson agrees the shares his TCP aimed for were wrong & supports our new direction
 - rewrite of IETF capacity sharing architecture in process
- the invisible hand of the market
 - favours ISPs that share capacity in their customers' best interests
 - based on theory of Hal Varian, now Chief Economist, Google
 - made practical by my team: congestion limiting within a flat fee
 - need to tweak TCP & IP (no change required to IP forwarding)

how to share an access cloud?

- once TCP/IP protocols can share internetwork capacity properly
 - partitioning access separately will be counter-productive


how Internet sharing 'works' endemic congestion & voluntary restraint




- those who take most, get most
 - voluntarily polite algorithm in endpoints
 - 'TCP-friendliness':



- a game of chicken – taking all and holding your ground pays



- or start more 'TCP-friendly' flows than anyone else (Web: x2, p2p: x5-100)





- or for much longer than anyone else (p2p file-sharing x200)
- net effect of both (p2p: x1,000-20,000 higher traffic intensity)

who is the fairest of them all?

1. equal bottleneck flow rates (TCP)?
2. access rate shared between active users, but weighted by fee (WFQ)?
3. volume caps tiered by fee?
4. heaviest applications of heaviest users throttled at peak times by deep packet inspection (DPI)?

none of the above harness end-system flexibility



- light usage can go much faster
- hardly affects completion time of heavy usage

NOTE: weighted sharing doesn't imply differentiated network service

- just weighted aggressiveness of end-system's rate response to congestion

a new resource accountability metric – a bandwidth trading unit




- how to measure
 - volume that is marked with explicit congestion notification (ECN)
 - can't be gamed by strategising machines
- a resource accountability metric
 - of customers to ISPs (too much traffic)
 - and ISPs to customers (too little capacity)
- a) cost to other users of your traffic
- b) marginal cost of equipment upgrade
 - so it wouldn't have been congested
 - so traffic wouldn't have affected others
- competitive market matches a) & b)

- cost of network usage
 - unforgivable for a business not to understand its costs
- answer: **congestion-volume**
 - volume weighted by congestion when it was sent
- takes into account all three factors

• bit-rate	✓	✓	✓	✓	✓
• weighted by congestion	✓	~	~	✗	~
• activity over time	✓	✗	✗	✓	✓

congestion-volume | TCP | WFQ | Vol | DPI



flat fee congestion policing



if ingress net could see congestion...

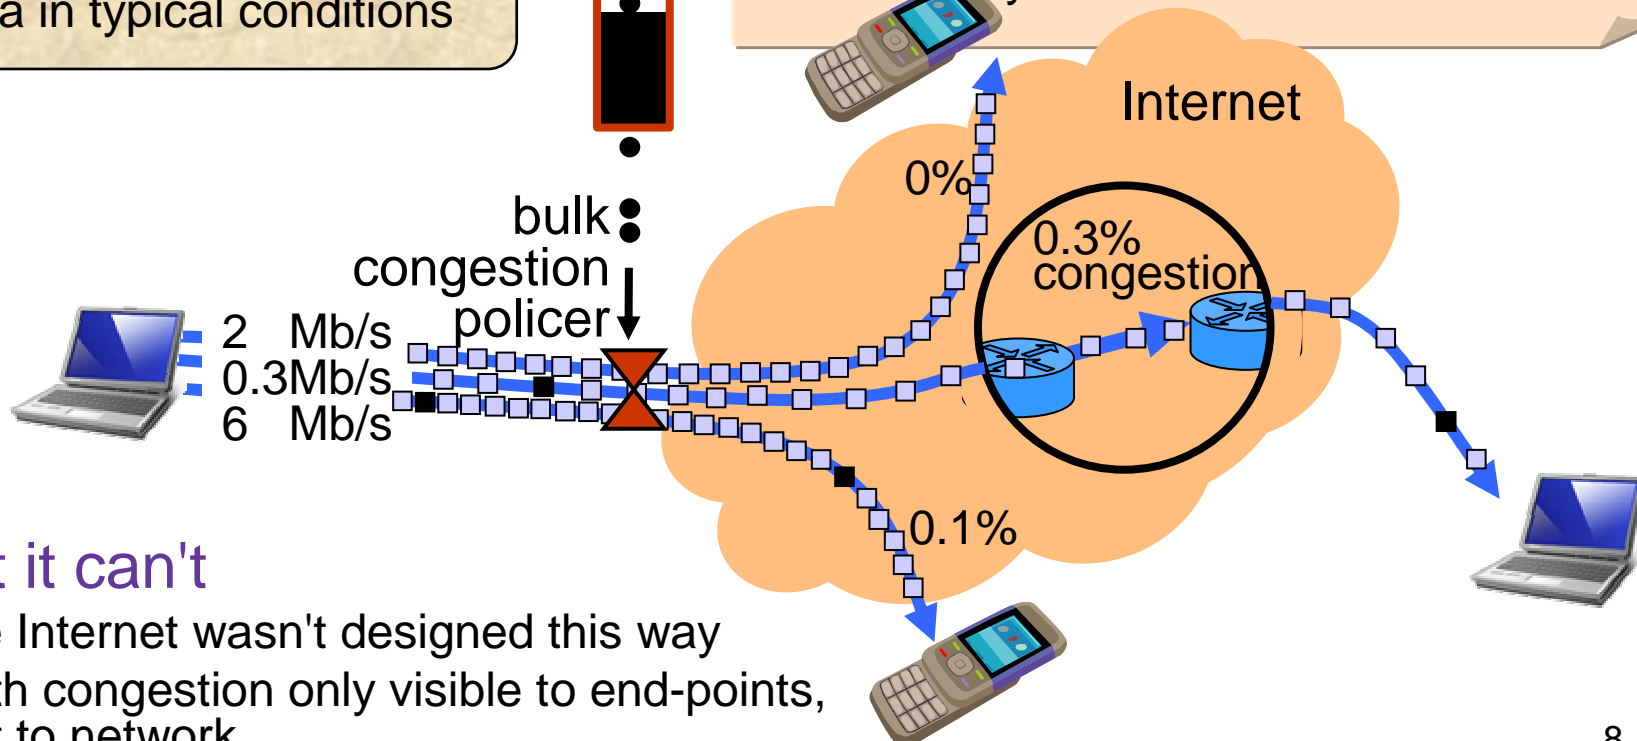
Acceptable Use Policy

'congestion-volume'
allowance: 1GB/month

@ £15/month

Allows ~70GB per day of
data in typical conditions


- incentive to avoid congestion
- simple invisible QoS mechanism
 - apps that need more, just go faster
- side-effect: stops denial of service
- only throttles traffic when your contribution to congestion in the cloud exceeds your allowance




...but it can't

- the Internet wasn't designed this way
- path congestion only visible to end-points, not to network

guaranteed bit-rate?
 or much faster 99.9% of the time?
 harnessing flexibility



- the idea that humans want to buy a known fixed bit-rate
 - comes from the needs of media delivery technology
 - hardly ever a human need or desire
- services want freedom & flexibility
 - access to a large shared pool, not a pipe
- when freedoms collide, congestion results
 - many services can adapt to congestion
 - shift around resource pool in time/space



% figures =
 no. of videos
 that fit into the
 same capacity

Constant Bit Rate **100%** Constant Quality **125%** Equitable Quality **216%**
 sequences encoded at same average of 500kb/s [Crabtree09]



closing off the future

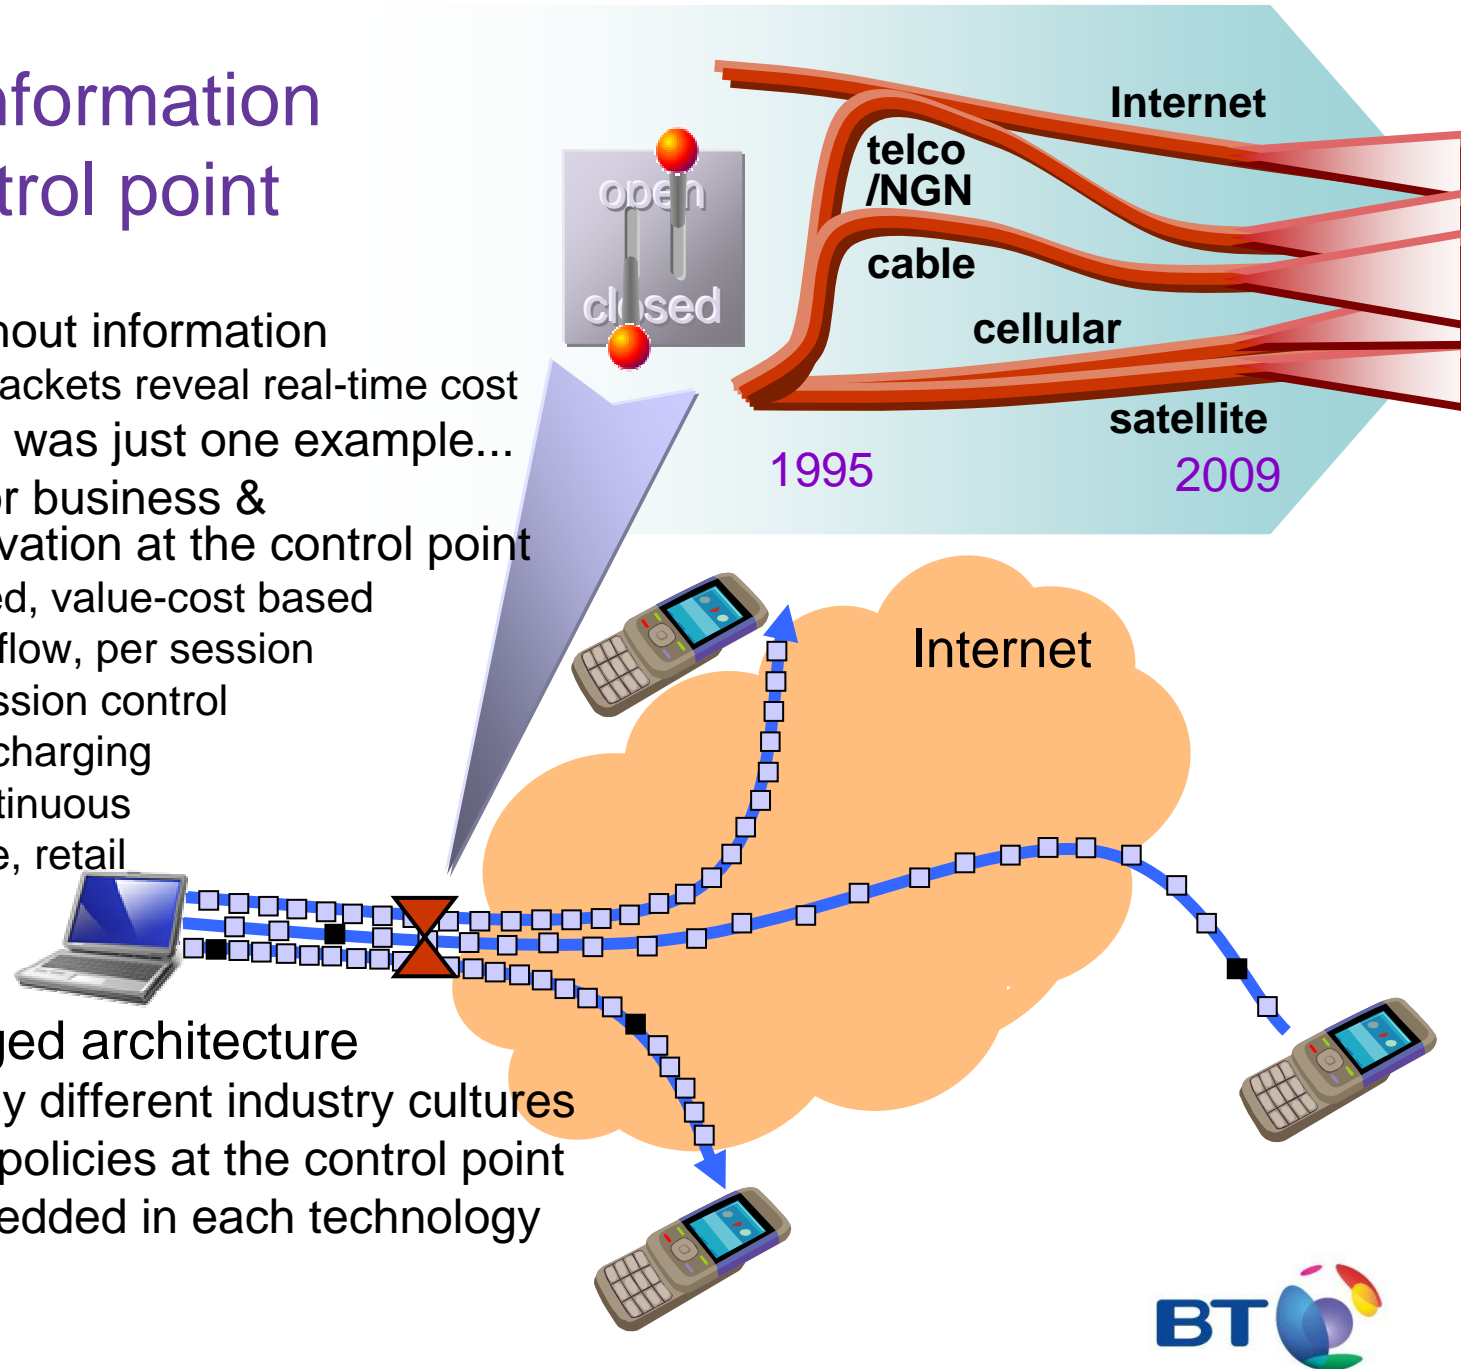
- ISPs must have a role in bandwidth sharing
 - minimally, incentivise end-systems to manage congestion
 - can't today, because ISPs can't see path congestion
- without correct metric, ISPs resort to application analysis
 - getting impossible to deploy a new use of the Internet
 - must negotiate the arbitrary blocks and throttles en route
- two confusable motives
 - fairer cost sharing
 - competitive advantage to own services
- how to deconfuse: make cost of usage transparent
 - fixing Internet technology should avoid need for legislation

bringing information to the control point

- no control without information
 - re-ECN packets reveal real-time cost
- flat fee policer was just one example...
- huge space for business & technical innovation at the control point

- cost based, value-cost based
- bulk, per flow, per session
- call admission control
- policing, charging
- tiers, continuous
- wholesale, retail

- truly converged architecture
 - can apply different industry cultures
 - through policies at the control point
 - not embedded in each technology





more info...

- The whole story in 5 pages
 - Bob Briscoe, "[A Fairer, Faster Internet Protocol](#)", IEEE Spectrum (Dec 2008)
- Inevitability of policing
 - [CFP06] The Broadband Incentives Problem, Broadband Working Group, MIT, BT, Cisco, Comcast, Deutsche Telekom / T-Mobile, France Telecom, Intel, Motorola, Nokia, Nortel (May '05 & follow-up Jul '06) <cfp.mit.edu>
- Slaying myths about fair sharing of capacity
 - [Briscoe07] Bob Briscoe, "[Flow Rate Fairness: Dismantling a Religion](#)" ACM Computer Communications Review 37(2) 63-74 (Apr 2007)
- How wrong Internet capacity sharing is and why it's causing an arms race
 - Bob Briscoe et al, "[Problem Statement: Transport Protocols Don't Have To Do Fairness](#)", IETF Internet Draft (Jul 2008)
- Understanding why QoS interconnect is better understood as a congestion issue
 - Bob Briscoe and Steve Rudkin "[Commercial Models for IP Quality of Service Interconnect](#)" BT Technology Journal 23 (2) pp. 171--195 (April, 2005)
- Equitable quality video streaming
 - [Crabtree09] B. Crabtree, M. Nilsson, P. Mulroy and S. Appleby "Equitable quality video streaming" Computer Communications and Networking Conference, Las Vegas, (January 2009)

Re-architecting the Internet:

The [Trilogy](http://www.trilogy-project.org) project <www.trilogy-project.org>

re-ECN & re-feedback project page:

<http://www.cs.ucl.ac.uk/staff/B.Briscoe/projects/refb/>

bob.briscoe@bt.com

the speed of sharing stretching Internet access

discuss...



main steps to deploy re-feedback / re-ECN



summary

rather than control sharing in the access links,
pass congestion info & control upwards

- network
 - turn on explicit congestion notification in data forwarding
 - already standardised in IP & MPLS
 - standards required for meshed network technologies at layer 2
(ECN in IP sufficient for point to point links)
 - deploy simple active policing functions at customer interfaces around participating networks
 - passive metering functions at inter-domain borders
- terminal devices
 - (minor) addition to TCP/IP stack of sending device
 - or sender proxy in network
- then new phase of Internet evolution can start
 - customer contracts & interconnect contracts
 - endpoint applications and transports

routing money

and simple internalisation of all externalities

