

deconfusing Internet traffic microeconomics

Bob Briscoe
Chief Researcher, BT Group
Oct 2008



why are we here?

- to understand, in order to set direction
- we could write a Wikipedia page (or FAQ) on Internet microeconomics

why are we confused?


- competition hasn't bitten
 - any half decent consultant's advice will lead to success
 - any half decent technology will do
- how things *are* working is poor guide to how they *will*
- has led to basic confusions even between value and cost
- entrance exam
 1. is charging by connectivity related to value or cost?
 2. is volume transferred a measure of value or of cost?
 3. does flat fee charging imply no variable costs?

	fixed	variable
value	connectivity	session type & volume
cost	dedicated access	how to share the cost of a cloud?




QoS: value \neq cost


UK Communications Service Revenue



UK Premium service revenues



UK Premium Service Volumes



- 'premium' = QoS demanding services
 - real-time VPN (e.g. corporate voice & videoconferencing)
 - premium BB services including broadband voice, fixed-mobile convergent services, video-telephony and IPTV/VoD
 - mobile voice (i.e. excluding SMS and MMS)
 - PSTN
- not necessarily *using* network QoS mechanisms (e.g. VoIP)
- 50% of premium revenues will depend on interconnect


Sources: Analysys Research (2005) and S Rudkin, BT internal report (Oct 2005)



proposed thought experiment

assume intense competition

- revenue driven to cost
 - eventually ensures customers, not providers, get the surplus value
 - all social policy heads this way – conditions a regulator is trying to create




- prepares for the inevitable
 - for service & network operators and equipment vendors
 - how to survive commoditisation
 - for architects & designers
 - an architecture that didn't foresee intense competition will be violated
 - e.g. deep packet inspection violated the Internet architecture
- useful assumption for a cross-industry working group
 - each can find our own ways to make margins above cost



how should we share the cost of a cloud?

- tremendous idea
 - anyone can use any link anywhere on the Internet without asking
- who decides how big a share each gets?
 1. TCP
 2. Comcast
 3. The Oval Office




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

the Internet way (TCP)	operators (& users)
'flow rate equality'	'volume accounting'
per data flow	per customer
instantaneous	over time

for scale: ~10M lines ringed in red




how Internet sharing 'works' endemic congestion & voluntary restraint




- aka. those who take most, get most


- technical consensus until Nov '06 [Briscoe07]
voluntarily polite algorithm in endpoints – 'TCP-fairness':



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




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



- or for much much longer than anyone else (p2p file-sharing x200)
- net effect of both (p2p: x1,000-20,000 higher traffic intensity) [Briscoe08]



Flow-Rate Fairness takes no account of activity




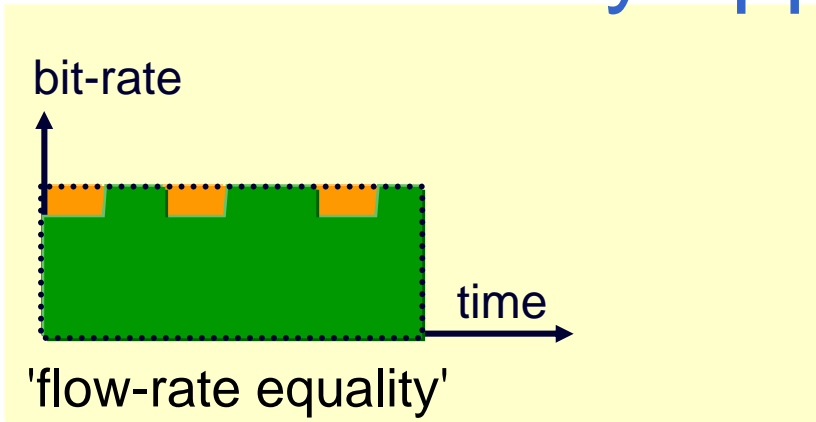
usage type	no. of users	activity factor	ave.simul flows /user	TCP bit rate /user	vol/day (16hr) /user	traffic intensity /user
attended	80	5%	=	417kbps	150MB	21kbps
unattended	20	100%	=	417kbps	3000MB	417kbps

x1

x20

x20

two arbitrary approaches fighting




the Internet way (TCP)


operators (& users)

degree of freedom	'flow rate equality'	'volume accounting'
multiple flows	x	✓
activity factor	x	✓
application control	✓	x
congestion variation	✓	x

- each cancels out the worst failings of the other
- Internet looks like 'it works OK'
- but the resulting arms race leaves collateral



fairer is faster – fixing the root cause



- enabler: limit congestion, not volume
- then end system congestion control will quickly evolve (cf. BitTorrent DNA)
 - heavy usage will back away whenever light usage appears
 - so **light** usage can go much faster
 - hardly affecting completion times of **heavy** usage
- differentiated QoS as *if* in the network



don't blame p2p

- commercial

Q. cost of network usage?

A. volume? NO

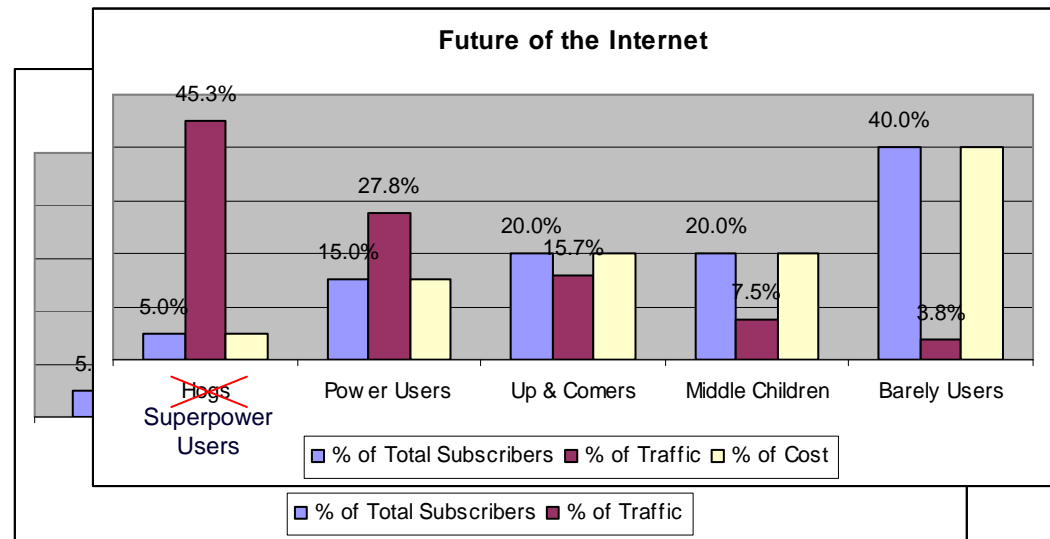
A. 'congestion volume'

- ISP's *own* unforgivable sloppiness over what their costs are

- technical

- lack of cost accountability in the Internet protocol (IP)

- machine-powered customers exploiting contracts & technology *ISPs* chose




who runs this pool?



not volume, but congestion volume: the missing metric

- not 'what you got' but 'what you unsuccessfully tried to get'
 - proportional to what you got
 - *but also* to congestion at the time
- 1. congestion volume: cost to other users
- 2. the marginal cost of upgrading equipment
 - so it wouldn't have been congested
 - so your behaviour wouldn't have affected others
- competitive market matches 1 & 2



*note: diagram is conceptual
congestion volume would be
accumulated over time
capital cost of equipment would be
depreciated over time*

NOTE: congestion volume isn't an extra cost

- part of the flat charge we already pay
- it's just the wrong people are paying it
- if we could measure who to blame for it we **might** see pricing like this...

access link	congestion volume allow'ce	charge
100Mbps	50MB/month	€15/month
100Mbps	100MB/month	€20/month

problems using congestion in contracts


	1. loss	2. ECN	3. re-ECN
can't justify selling an impairment	☹	☺	☺
absence of packets is not a contractible metric	☹	☺	☺
congestion is outside a customer's control	☹	☹	☺
customers don't like variable charges	☹	☹	☺
congestion is not an intuitive contractual metric	☹	☹	☹

1. **loss**: used to signal congestion since the Internet's inception

- computers detect congestion by detecting gaps in the sequence of packets
- computers can hide these gaps from the network with encryption

2. **explicit congestion notification [ECN]**: standardised into TCP/IP in 2001

- approaching congestion, a link marks an increasing fraction of packets
- implemented in Windows Vista (but off by default) and Linux, and IP routers (off by default)



3. **re-inserted ECN [re-ECN]**: standards proposal since 2005 (not formal IETF yet)

- packet delivery conditional on sender declaring expected congestion
- uses ECN equipment in the network unchanged

Acceptable Use Policy

Your 'congestion volume' allowance:
1GB/month (= 3kb/s continuous)

This only limits the traffic you can try to transfer above the maximum the Internet can take when it is congested.

Under typical conditions this will allow you to transfer about **70GB per day**.

If you use software that seeks out uncongested times and routes, you will be able to transfer a lot more.

Your bit-rate is otherwise unlimited

limiting congestion?

- only throttles traffic when contribution to congestion elsewhere exceeds allowance [Jacquet08]
- otherwise free to go at any bit-rate

congestion · bit-rate

$$0\% \cdot 2 \text{ Mb/s} = 0.0\text{kb/s}$$

$$0.3\% \cdot 0.3\text{Mb/s} = 0.9\text{kb/s}$$

$$0.1\% \cdot 6 \text{ Mb/s} = 6.0\text{kb/s}$$

$$6.9\text{kb/s}$$

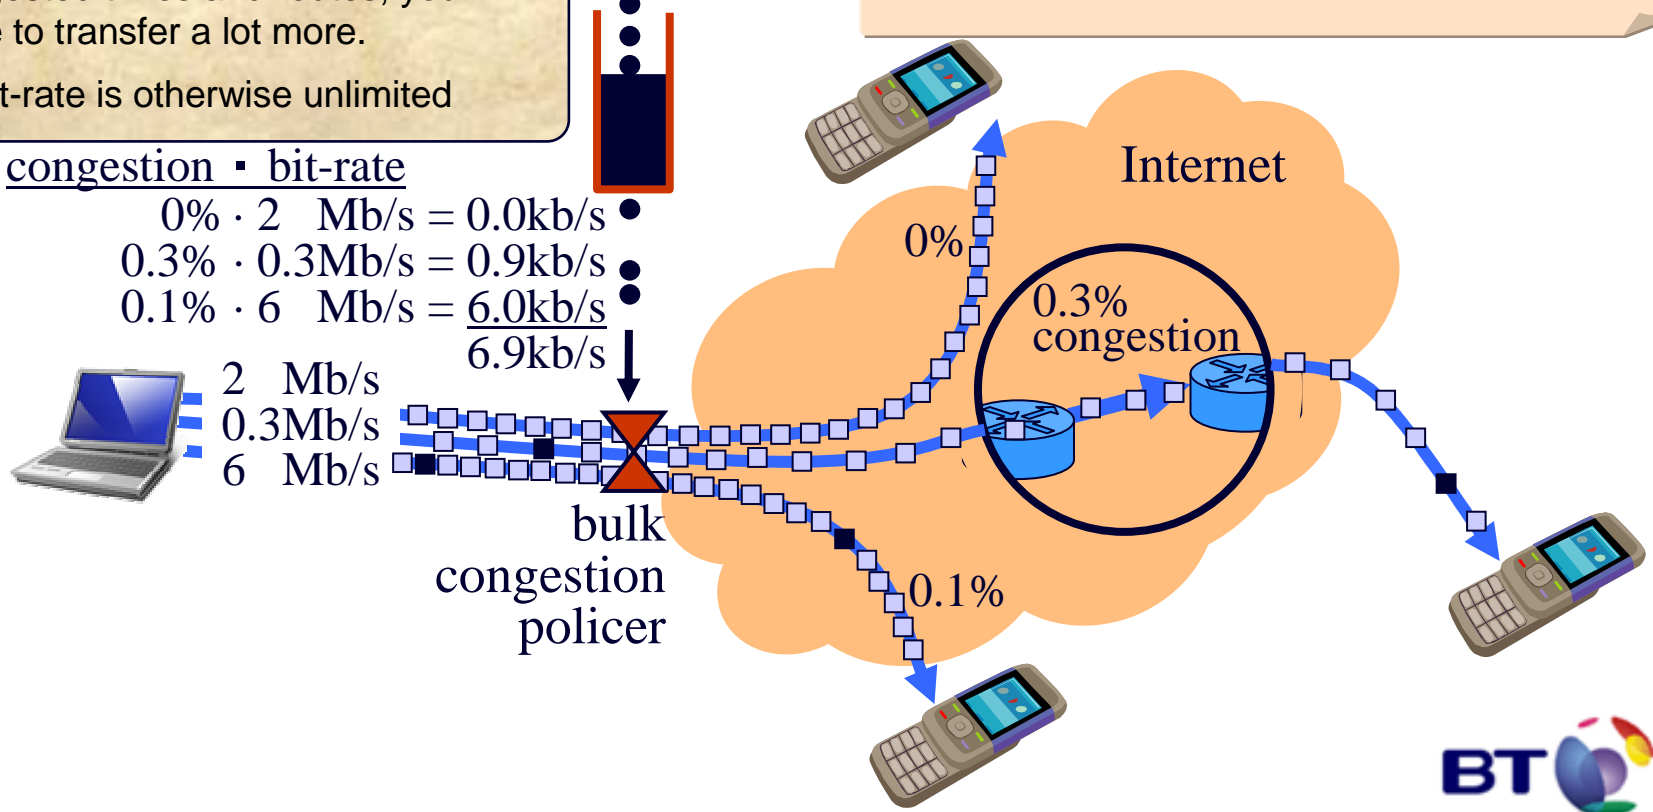


2 Mb/s


0.3Mb/s

6 Mb/s

bulk
congestion
policer



automatic interconnect usage cost allocation



interconnect aggregation


simple internalisation of all externalities
'routing money'

legend:

re-ECN
downstream
congestion
marking [%]

bit rate

area =
instantaneous
downstream
congestion
volume



0|0|2|7|6|0|5

just two counters at border,
one for each direction

meter monthly bulk volume
of packet markings

= aggregate downstream
congestion volume in flows
without measuring flows



summary

- assuming competition: deconfuses the economics
- for industry players the future becomes clearer
- the regulator's & architect's tasks become clearer

- there's still problems to unravel
 - semi-experts co-authoring a Wiki FAQ would help unpick them



high hanging fruit

- if sharing a cloud is sorted out in IP
- could remove bit-rate limits in shared access links?
 - remove multiple access from cable, wireless, PON?
- example
 - 100 users sharing a 10G PON
 - could all peak at 10G
 - not 100M
- caveat: scary e2e congestion control dynamics

more info...

- Inevitability of policing
 - [BBincent06] 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>
- Stats on p2p usage across 7 Japanese ISPs with high FTTH penetration
 - [Cho06] Kenjiro Cho et al, "The Impact and Implications of the Growth in Residential User-to-User Traffic", In Proc ACM SIGCOMM (Oct '06)
- 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
 - [Briscoe08] 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
 - [Briscoe05] Bob Briscoe and Steve Rudkin "[Commercial Models for IP Quality of Service Interconnect](#)" BT Technology Journal 23 (2) pp. 171--195 (April, 2005)
- Re-architecting the Future Internet:
 - The [Trilogy](#) project
- Re-ECN & re-feedback project page, includes [re-ECN, Jacquet08, Briscoe07, Briscoe08]:
<http://www.cs.ucl.ac.uk/staff/B.Briscoe/projects/refb/>
- These slides
<www.cs.ucl.ac.uk/staff/B.Briscoe/present.html>



further references

- [Clark05] David D Clark, John Wroclawski, Karen Sollins and Bob Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet," IEEE/ACM Transactions on Networking (ToN) 13(3) 462–475 (June 2005) <portal.acm.org/citation.cfm?id=1074049>
- [MacKieVarian95] MacKie-Mason, J. and H. Varian, "Pricing Congestible Network Resources," IEEE Journal on Selected Areas in Communications, `Advances in the Fundamentals of Networking' 13(7)1141--1149, 1995
<http://www.sims.berkeley.edu/~hal/Papers/pricing-congestible.pdf>
- [Kelly98] Frank P. Kelly, Aman K. Maulloo, and David K. H. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. Journal of the Operational Research Society, 49(3):237–252, 1998
- [Gibbens99] Richard J. Gibbens and Frank P. Kelly, Resource pricing and the evolution of congestion control, Automatica 35 (12) pp. 1969—1985, December 1999 (lighter version of [Kelly98])
- [Gibbens02] Richard J. Gibbens and Frank P. Kelly , "On Packet Marking at Priority Queues" In: IEEE Transactions on Automatic Control 47 (6) pp. 1016--1020 (June, 2002).
- [ECN] KK Ramakrishnan, Sally Floyd and David Black "The Addition of Explicit Congestion Notification (ECN) to IP" IETF RFC3168 (Sep 2001)
- [Briscoe05] Bob Briscoe, Arnaud Jacquet, Carla Di-Cairano Gilfedder, Andrea Soppera and Martin Koyabe, "Policing Congestion Response in an Inter-Network Using Re-Feedback" In: Proc. ACM SIGCOMM'05, Computer Communication Review 35 (4) (September, 2005)
- Policing Freedom





deconfusing the issues
Internet traffic economics



Q&A



capacity costs



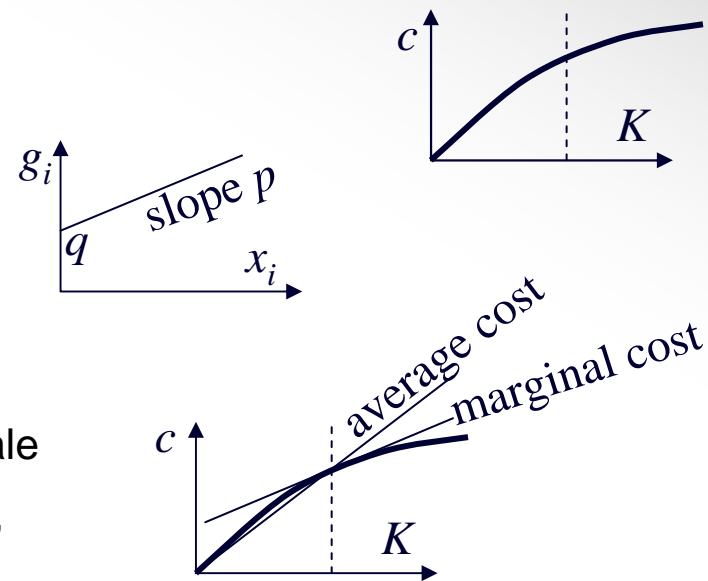
- selling QoS = managing risk of congestion
 - if no risk of congestion, can't sell QoS
 - congestion risk highest in access nets (cost economics of fan-out)
 - also *small* risk in cores/backbones (failures, anomalous demand)

usage vs subscription prices

Pricing Congestible Network Resources [MacKieVarian95]

- assume competitive providers buy capacity K [b/s] at cost rate [€/s] of $c(K)$
- assume they offer a dual tariff to customer i
 - subscription price q [€/s]
 - usage price p [€/b] for usage x_i [b/s], then charge rate [€/s], $g_i = q + px_i$
- what's the most competitive choice of p & q ?

- $\frac{\text{usage revenue}}{\text{capacity cost}} = \frac{1}{e}$ where e is elasticity of scale
 - if charge less for usage and more for subscription, quality will be worse than competitors
 - if charge more for usage and less for subscription, utilisation will be poorer than competitors



$$e = \frac{\text{average cost}}{\text{marginal cost}}$$

$$= \frac{c(K)}{K} \cdot \frac{1}{c'(K)}$$



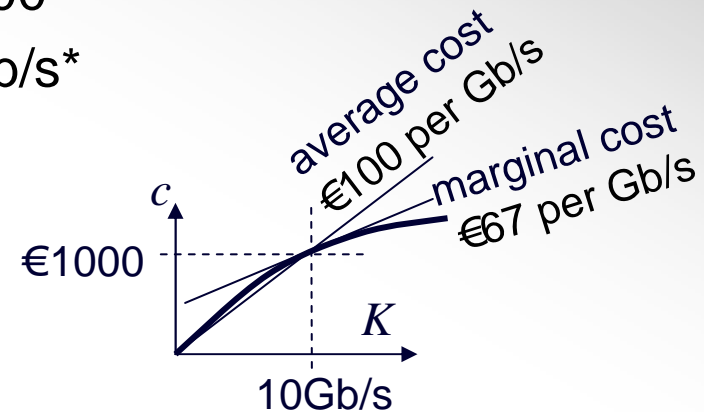
toy example

- if a 10Gb/s link interface costs €1000
- and it costs €67 to upgrade to 11Gb/s*
 - average cost = €100 per Gb/s
 - marginal cost ~ €67 per Gb/s

$$e = \frac{\text{average cost}}{\text{marginal cost}} = \frac{3}{2}$$

$$\therefore \frac{\text{usage revenue}}{\text{capacity cost}} = \frac{1}{e} = \frac{2}{3} \qquad \frac{\text{subscription revenue}}{\text{capacity cost}} = \frac{1}{3}$$


- ie usage revenue covers marginal cost
subscription revenue covers the rest
- then add operational costs



* obviously not practical to physically upgrade in such small steps

cost-shifting between services


- scenario
 - ISP also a higher level service provider (TV, video phone, etc)
 - competing with independent service providers (Skype, YouTube, etc)
- capacity & QoS costs for high value services
 - ISP buys capacity & QoS internally
 - independent SP can just take as much best-efforts bandwidth as they need
 - because of how Internet sharing 'works'
- cost of heavy usage service can be subsidised by ISP's lighter users




p2p quickly fills up fibre to the home

Distribution of customers' daily traffic into & out of a Japanese ISP (Feb 2005)

(5GB/day equivalent to
0.46Mbps if continuous)



Changing technology shares of Japanese access market



Courtesy of Kenjiro Cho et al
The Impact and Implications of the Growth
in Residential User-to-User Traffic, SIGCOMM (Oct '06)

