

the cost of freedom

Bob Briscoe
Chief Researcher, BT Group
and UCL
Dec 2006




degrading specific Internet applications

- a trend with two possible causes
 - technical deficiency : Internet market : global :
 - regulatory deficiency : access market : mostly US-specific
- outline of talk
 1. technical deficiency with current Internet
 - resource allocation architecture allows those who bully most to get most (p2p, video)
 - network cannot arbitrate, because key usage information inaccessible to it
 - lacking a proper solution, operators kludge it by degrading likely culprit apps
 - intentions can be honourable, but confusable with degrading their competition
 - confusion can be conveniently exploited by either side
 2. proposed solution to both these problems (and more)
 - loose global grouping of researchers re-architecting Internet since 2000
 - 1-bit app-neutral fix to the Internet Protocol, in early standards process
- purpose of talk
 - does the proposed solution create a playing field you would be happy with?
 - if so, it needs support - cannot sustain through standards without a lot more help

“freedom to use my line how I want” limits freedom of others

- no. of access lines that can congest any other Internet link
 - staying around 1,000 – 100,000
 - cf. ~10M lines ringed in red
- congestion on shared links
 - not just a technical issue
 - continual conflict betw. real people & between real businesses

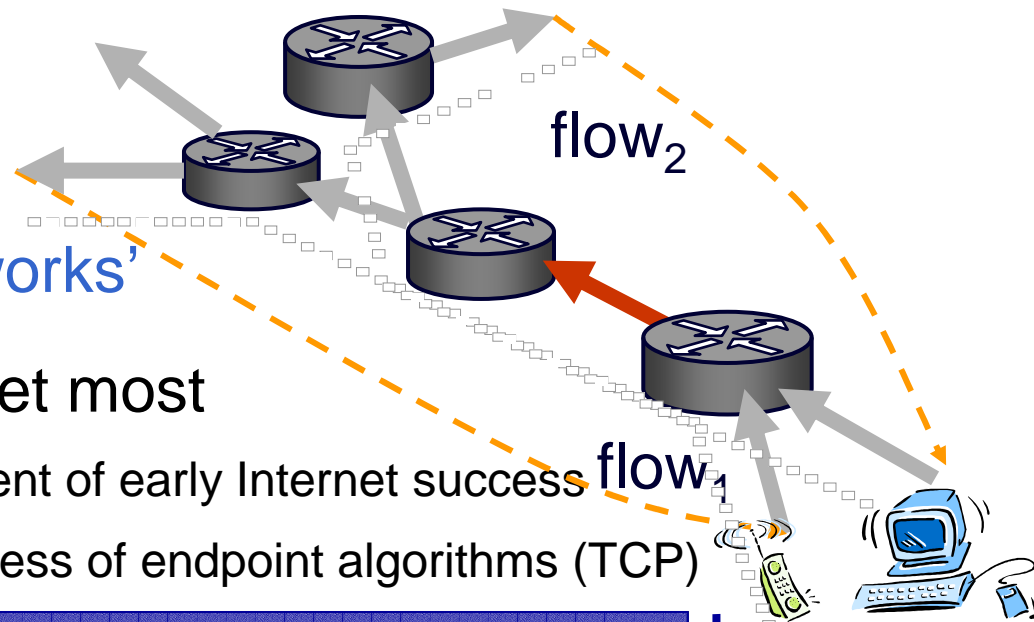


Internet topology visualization produced by Walrus
(Courtesy of Young Hyun, CAIDA)


- need better usage control not just access link capacity
 - but volume is an inadequate metric
 - so ISPs resort to targeting apps




freedom how Internet sharing 'works'




- those who push most, get most
 - restraint: the other ingredient of early Internet success
 - reliant on voluntary politeness of endpoint algorithms (TCP)



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



- or starting more 'TCP-fair' flows than anyone else



- or for much longer than anyone else (p2p file-sharing)



ineffective kludges are making matters worse fuelling adversarial climate


- deep packet inspection (DPI) cannot win arms race against obfuscation
 - 80% of payloads now carry randomised app identifier
 - latest p2p apps use payload encryption & imitate other apps
 - more & more false positives, more customer support calls

download
bit rate



customer of an ISP using DPI to throttle p2p turns *off* encryption in BitTorrent client

- intentions might be honourable
 - protecting the many from the few
- but counter-productive
 - if easily bypassed and easily turned against itself
 - if (mis)interpretable as discriminating against competition



← 200kbps
(2M contended)

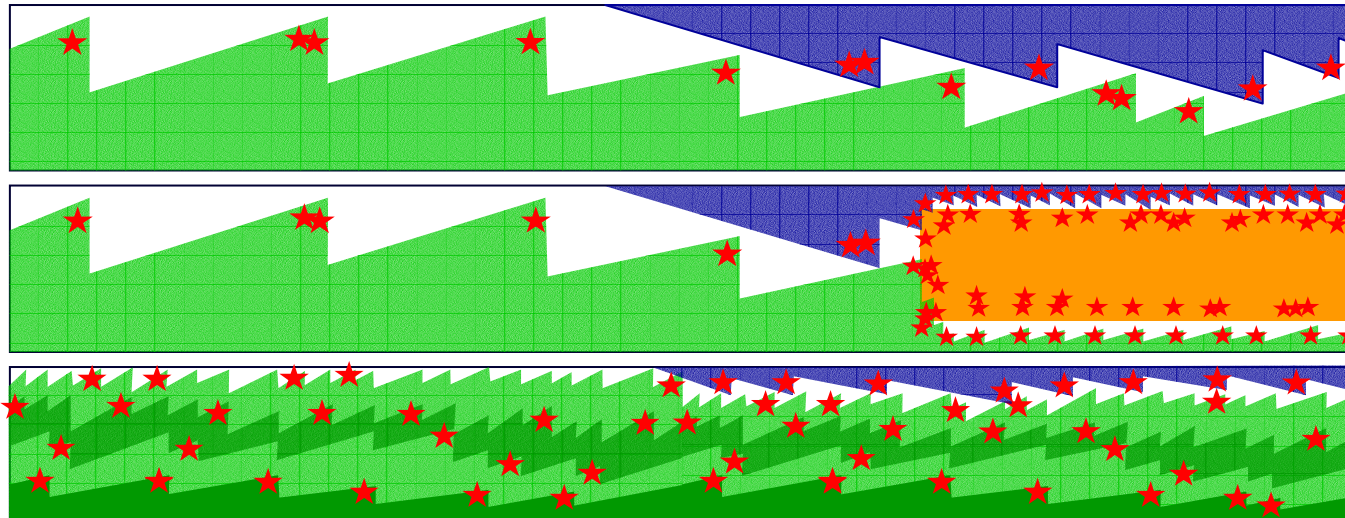
← 30kbps



the missing metric

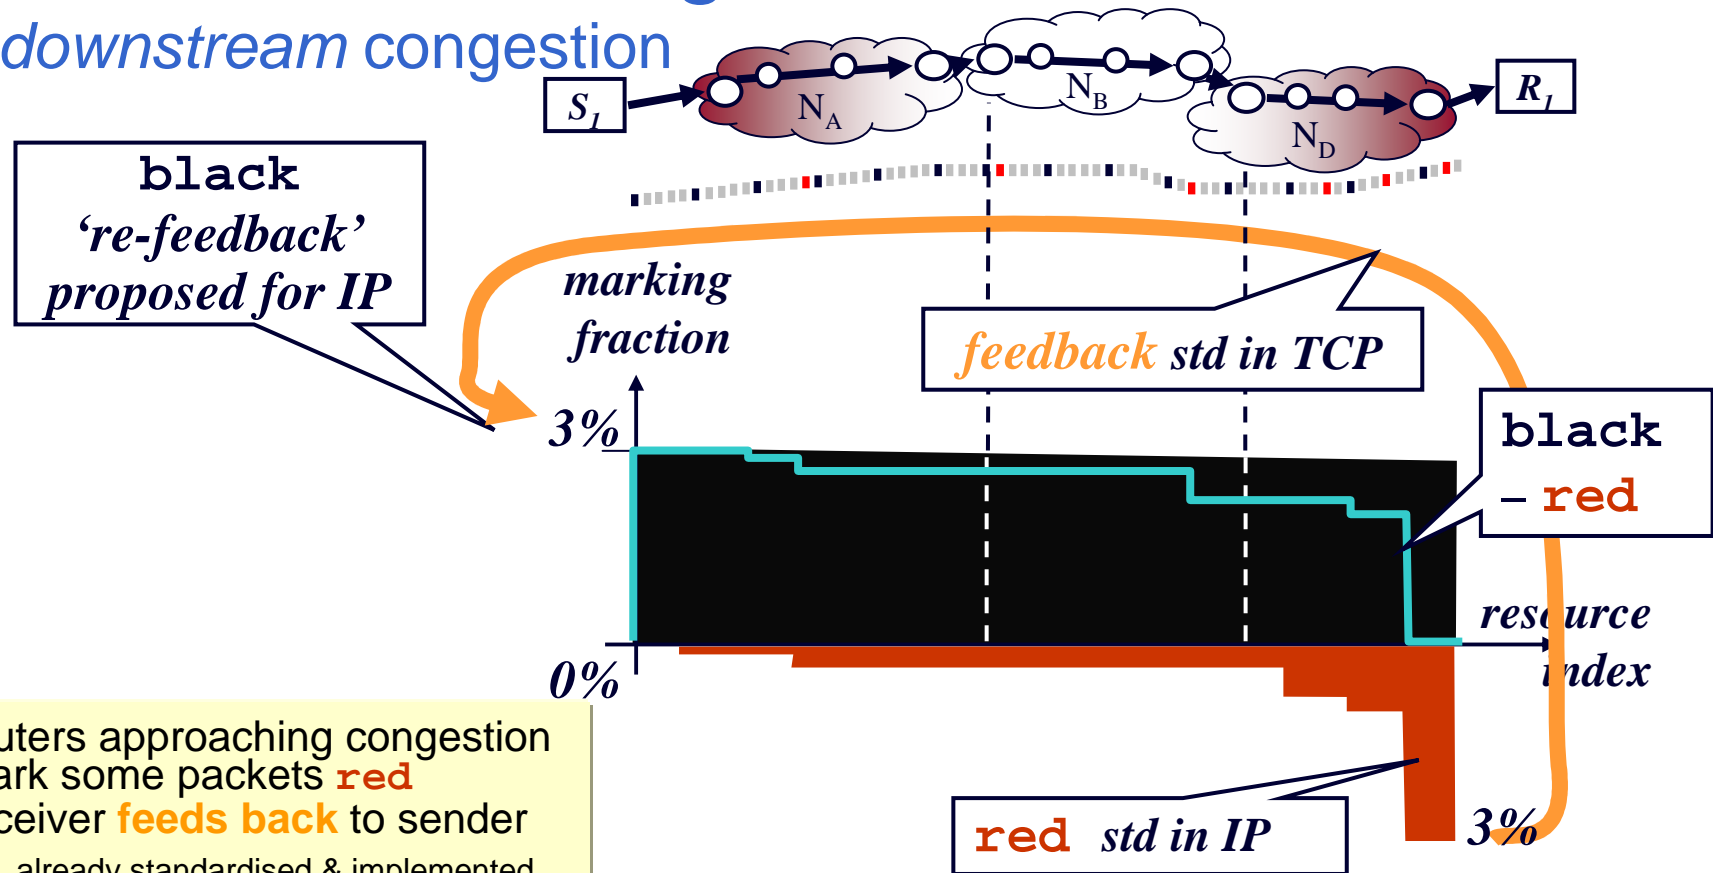
congestion volume (not volume)

- a competitive market tends towards cost
 - cost of each customer's usage is "stuff sent that can't get through"
 - but impractical to measure absence of stuff
 - explicit congestion notification (ECN): standardised into IP in 2001
 - mark "stuff that wouldn't get through" if congestion got worse



the other missing metric

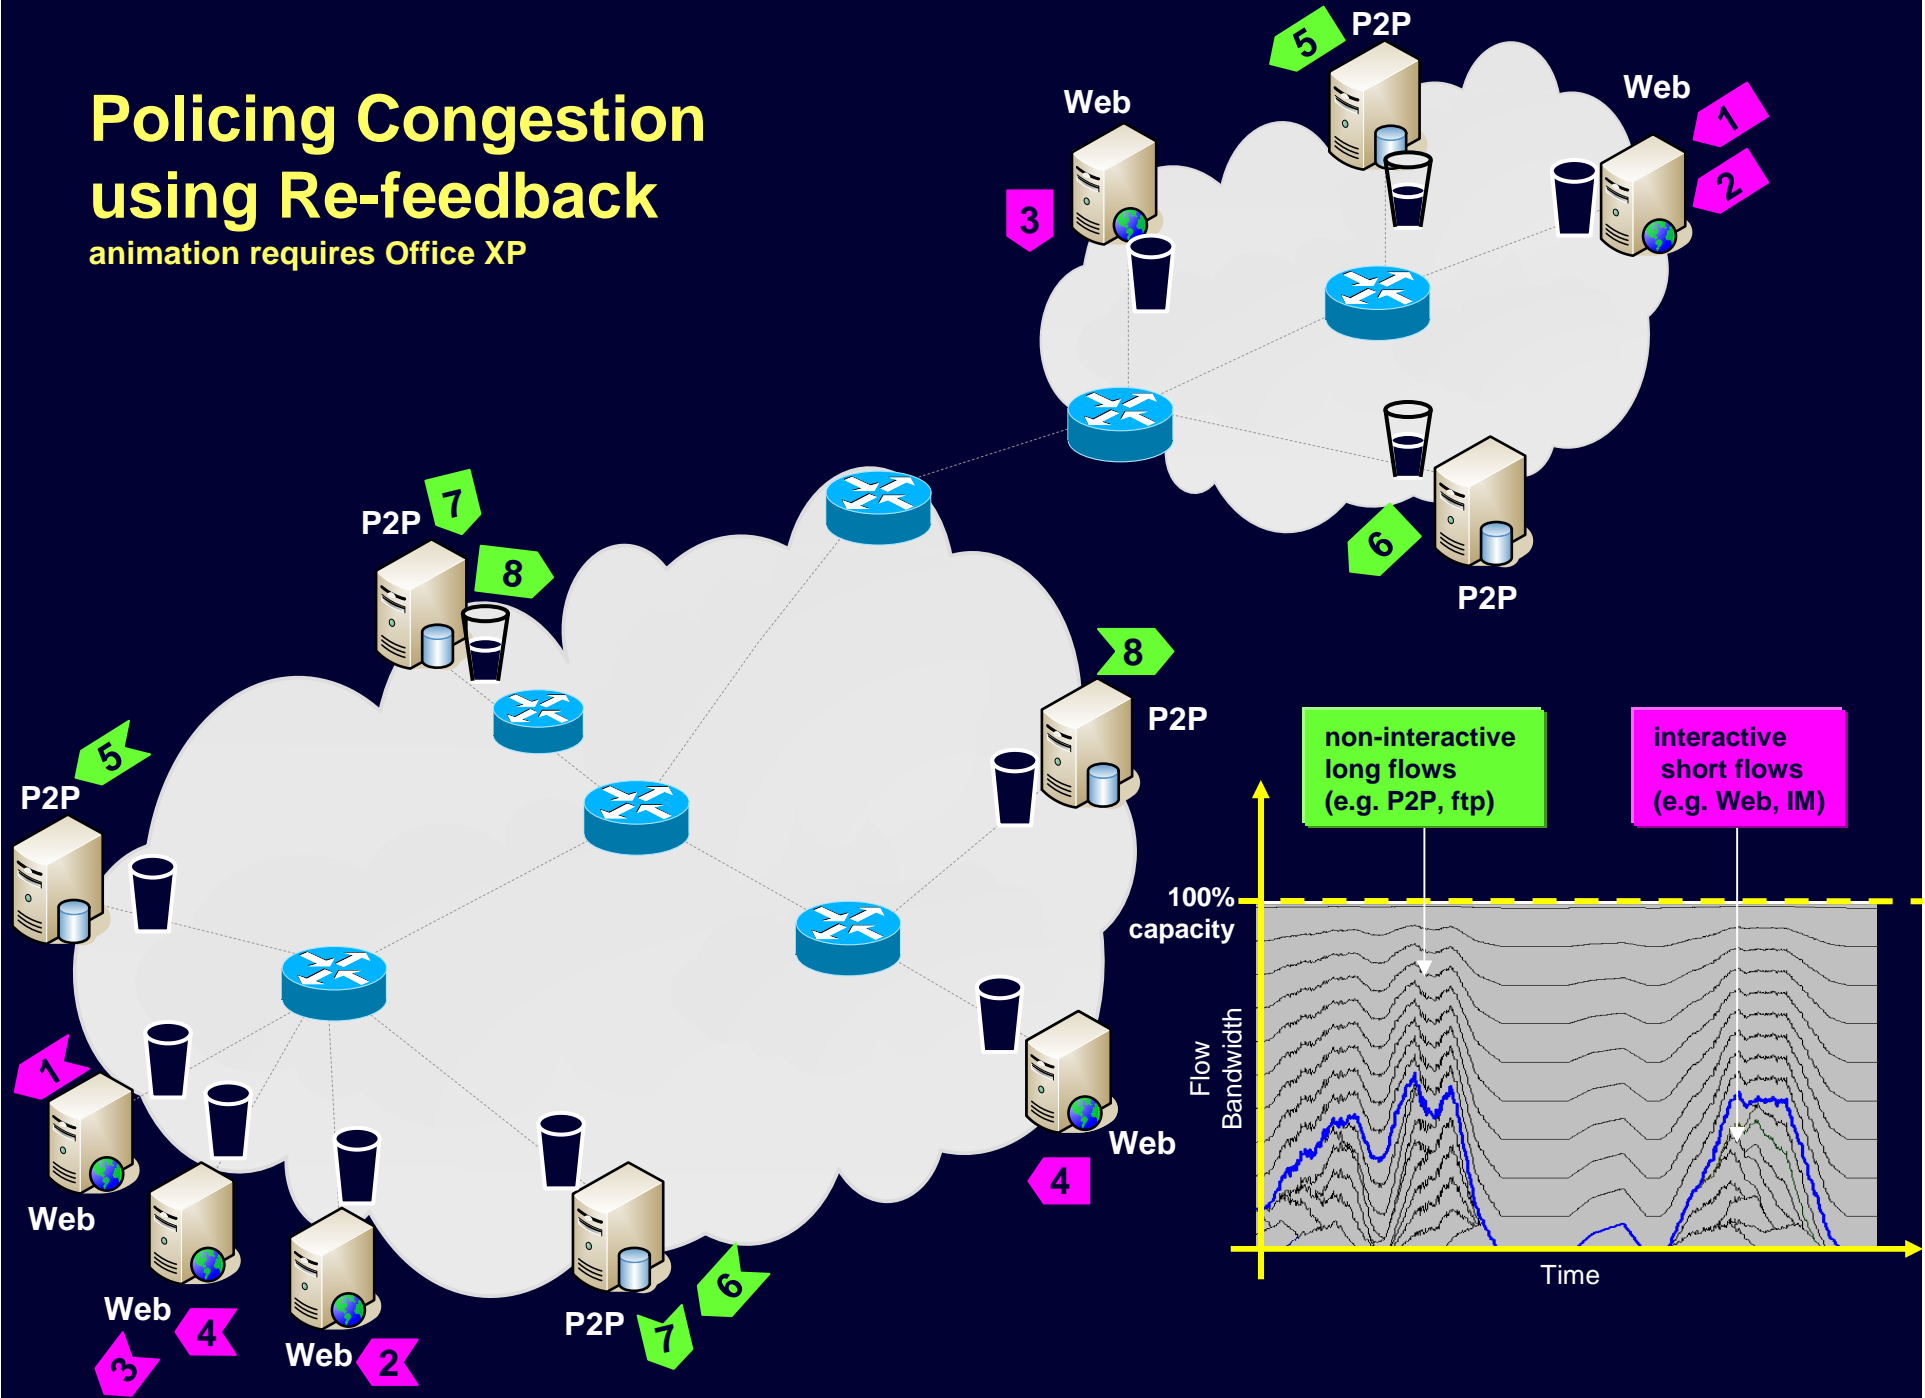
downstream congestion



- routers approaching congestion mark some packets **red**
- receiver **feeds back** to sender
 - already standardised & implemented
 - not generally turned on by operators
- sender re-inserts **feedback** by marking packets **black**
 - **re-feedback** requires standardisation


Policing Congestion using Re-feedback

animation requires Office XP



using the downstream congestion metric

one example: per-user **policer**



two different customers, same deal

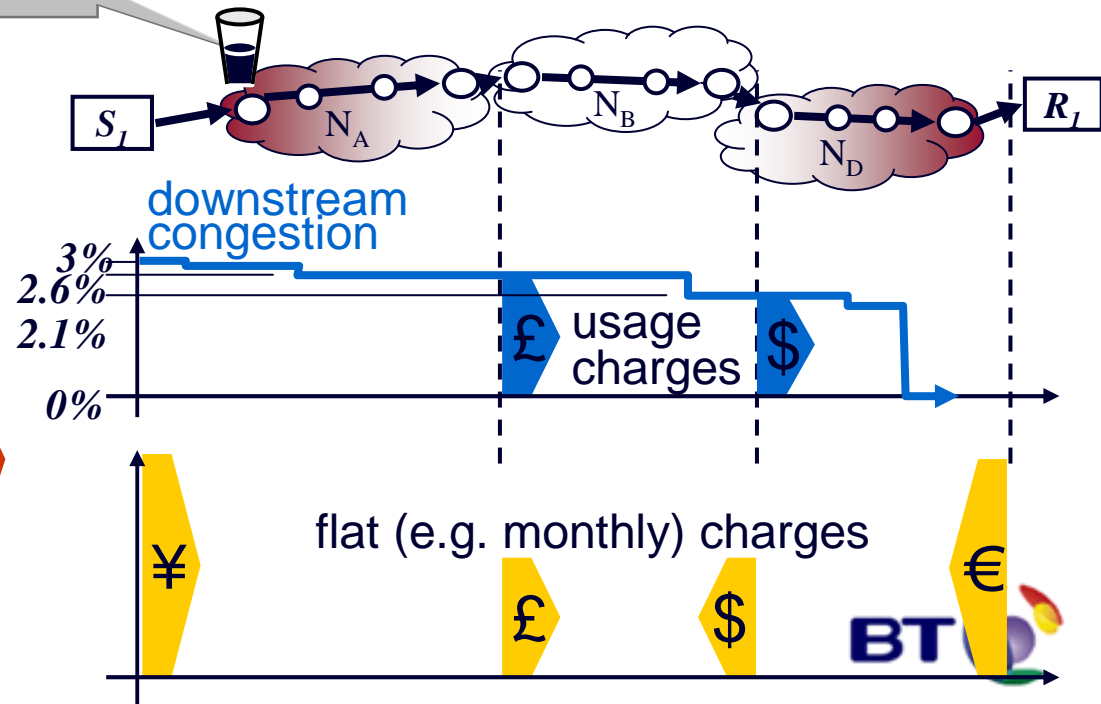

interactive short flows
(e.g. Web, IM)

other examples

- make flows respond to congestion (VoIP, video, DDoS)
- no policing at all

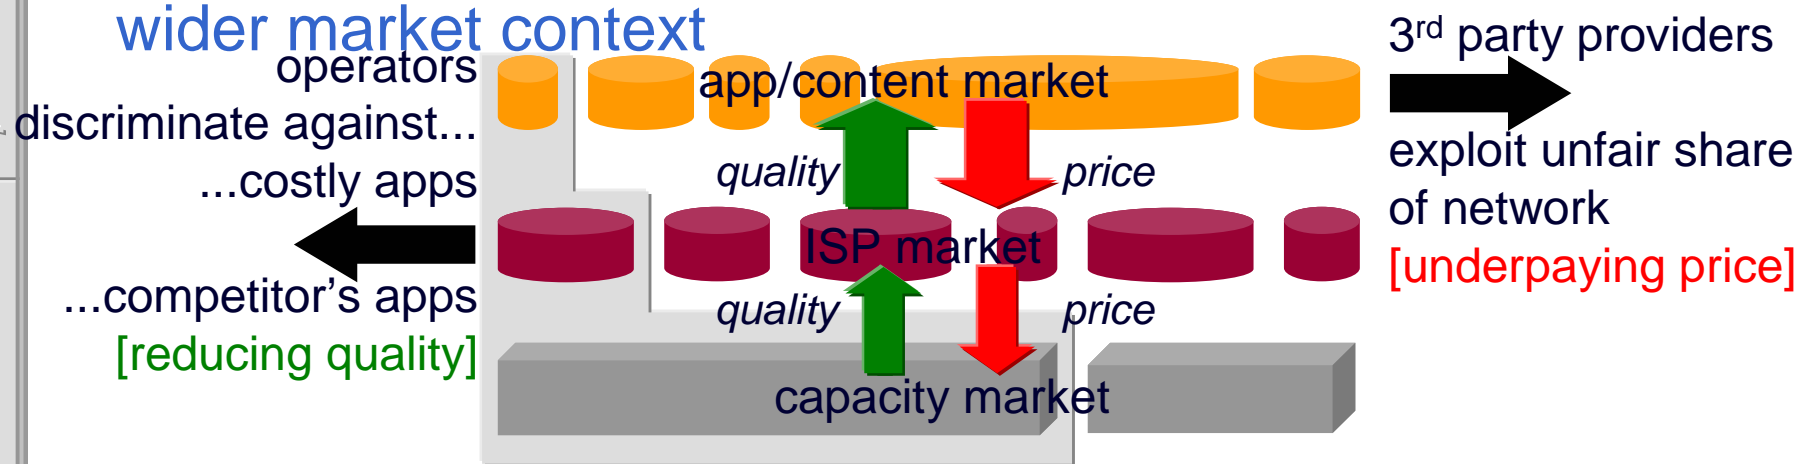
inter-domain accountability for congestion

- metric for inter-domain SLAs or usage charges
 - N_B applies penalty to N_A in proportion to bulk volume of **black** less bulk volume of **red** over, say, a month
 - could be tiered penalties, directly proportionate usage charge, etc.
 - flows de-aggregate precisely to responsible networks
 - N_A can deploy **policer** to prevent S_1 costing more than revenue



degrading specific Internet applications

wider market context



- solution: identify costly *bits*
 - then quality can rise to match willingness to pay

| market | problem | appropriate remedy | inappropriate remedy |
|----------|-----------------------|--------------------------|------------------------------|
| Internet | architecture | fix architecture | US net neutrality regulation |
| access | weak competition (US) | fix US access regulation | US net neutrality regulation |
| | going well (e.g. UK) | no change | |

designed for tussle

Internet needs all these answers – market selection finds

balance

demand side – freedom to degrade others

- the Internet is all about the freedom to get what I want (within my line rate)

limited by how much I impinge on the freedom of others

- enforceable congestion control

freedom within fairness

- differentiated quality of service

you'll get what you contract to have


- you'll get what we infer you want given what you're doing

supply side – freedom to degrade competitors

architecture allows extremes but doesn't help them and provides handles for the market to make it very hard for them

summary

- Internet needs to be able to discriminate
 - against bits limiting the freedom of others – *bits* causing congestion
 - then wouldn't need to discriminate against *apps* causing congestion
- operators can choose not to limit their users' freedoms
 - but they take responsibility for congestion their users cause in other nets
- if operators do discriminate against apps
 - customers need enough choices to be able to switch operators
 - or apps can often obfuscate themselves anyway
- these economic effects require change to the Internet Protocol
 - making IP more suitable as the basis of a converged architecture
 - early in standards process – link on next slide
 - please assess it – urgently needs support through standards



the cost of freedom

[<www.cs.ucl.ac.uk/staff/B.Briscoe/present.html#0612crn>](http://www.cs.ucl.ac.uk/staff/B.Briscoe/present.html#0612crn)

Q&A


& spare slides




capacity growth will prevent congestion?

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

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



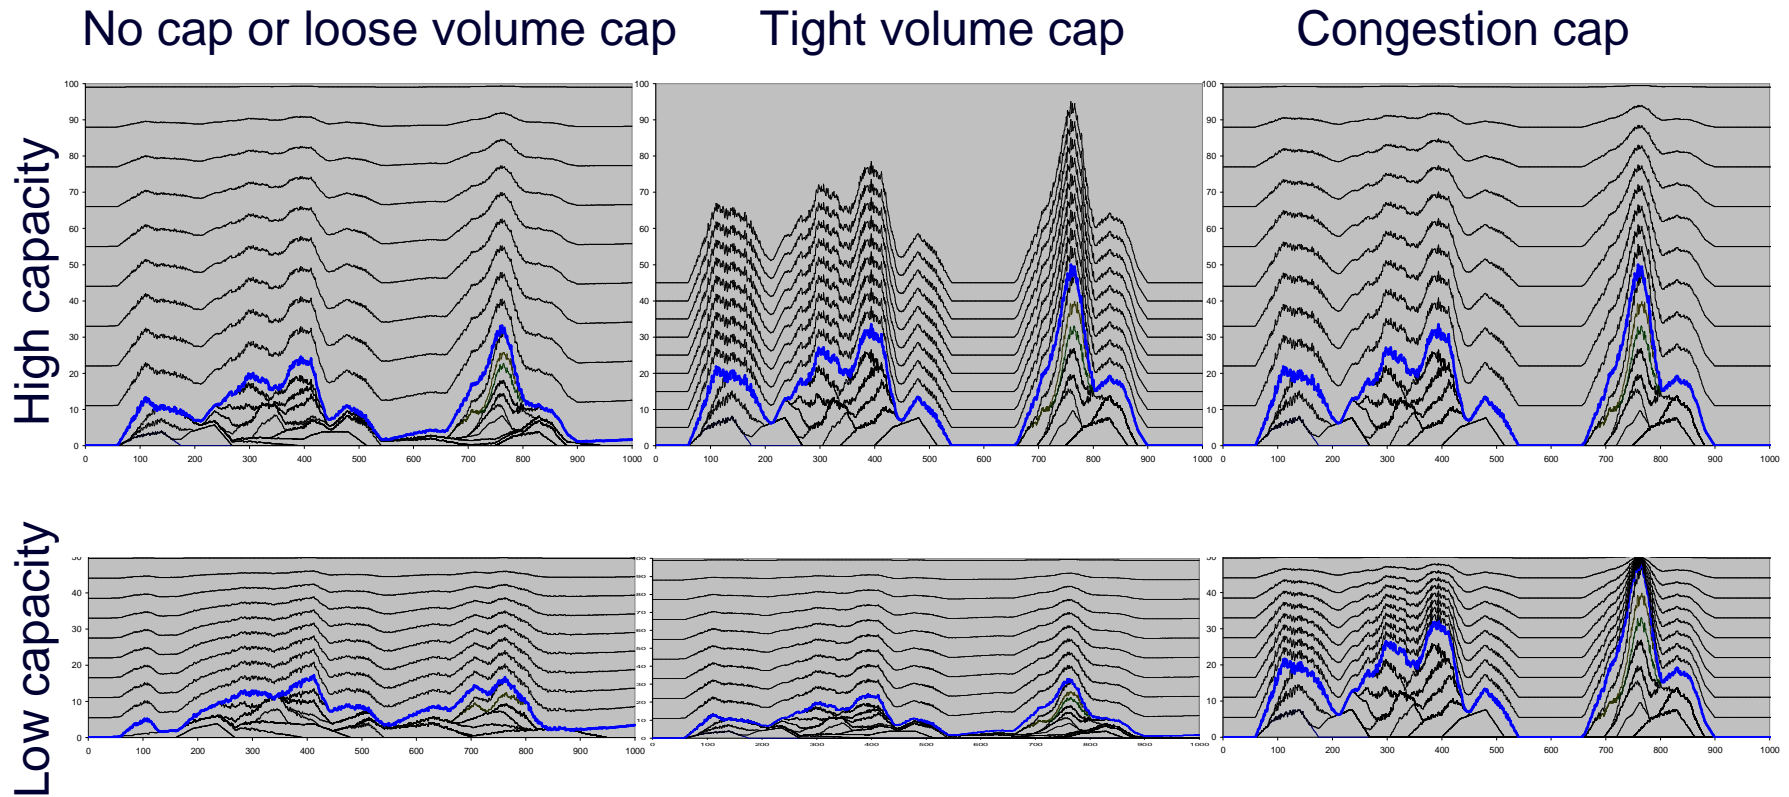
Changing technology shares of Japanese access market



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




congestion cap auto-adjusts volume cap always a hard compromise



incentives


solution step #4



- won't sender or receiver simply understate congestion?
- no – drop enough traffic to make fraction of **red** = **black**
- goodput best if rcvr & sender honest about feedback & re-feedback

aggregation

internalisation of externalities



prime directive for re-architecting the Internet

- ‘design for tussle’
 - design so major social & economic struggles can be resolved at ‘run-time’
 - by market or social regulation
 - if instead you try to embed political outcomes into the architecture at ‘design-time’
 - powerful economic forces will ignore/violate the architecture to achieve their ends
- Internet > public Internet
 - eg. in BT’s 21C network, public Internet & PSTN replacement are just IP virtual private networks
- resource allocation architecture
 - but current Internet inherently open (to abuse)
 - re-feedback allows more control
 - making IP more generic as the basis of a converged architecture
 - then each operator’s choice between open & closed can be determined by market forces
 - but the whole spectrum of choices can be neutral to specific apps

