

Internet resource sharing: a way forward?

Bob Briscoe
Chief Researcher, BT
May 2009

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




fairness



- one would expect ISPs to care about fairness
 - ISPs with poor fairness will lose customers to competitors
 - ISPs never cared about fairness between flow rates
- flow rate fairness: invention of protocol community
 - completely unrelated to fairness in real life
 - myopically looks at each flow separately, not customers
 - myopically looks at each instant, not over time
- ISPs use volume/month as a fairness metric
 - it counts across flows
 - and over time
 - ...

TCP-friendly meaningless over time

time is unfortunately real



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

target structure: *network* fairness




⇒ ⊗ × bottleneck policers: active research area since 1999

- detect flows causing unequal share of congestion
- located at each potentially congested router
- takes no account of how active a source is over time
- nor how many other routers the user is congesting
- based on cheap pseudonyms (flow IDs)



✓ re-ECN / ECN


- like counting volume, but 'congestion-volume'
- reveals congestion caused in all Internet resources by all sources (or all sinks) behind a physical interface, irrespective of addressing
- accumulates over time
- no advantage to split IDs
- focus of fairness moves from flows to packets



core of solution

congestion-volume metric

- congestion-volume
 - your volume weighted by link congestion when each packet is served
- intuition
 - some ISPs count volume only during peak
 - like counting (100% x volume) during peak and (0% x volume) otherwise
 - congestion-volume $C \equiv \int p(t)x_i(t) dt$
 - cf. straight volume $V \equiv \int x_i(t) dt$
- measurement
 - the amount of data discarded from your traffic
 - or marked with explicit congestion notification (ECN)
 - end-point function in current architecture



1. cost to other users of your traffic
 2. the marginal cost of upgrading equipment
 - so it wouldn't have been congested
 - so traffic wouldn't have affected others
- competitive market matches 1 & 2

metric for customers to judge ISPs,
and ISPs to judge customers

congestion = too much traffic meets too little capacity

most interesting when 'congestion' = marking, not loss

*note: diagram is conceptual
congestion volume & equipment capex
would be accumulated over time*



a vision: flat fee congestion policing



if ingress net could see congestion...

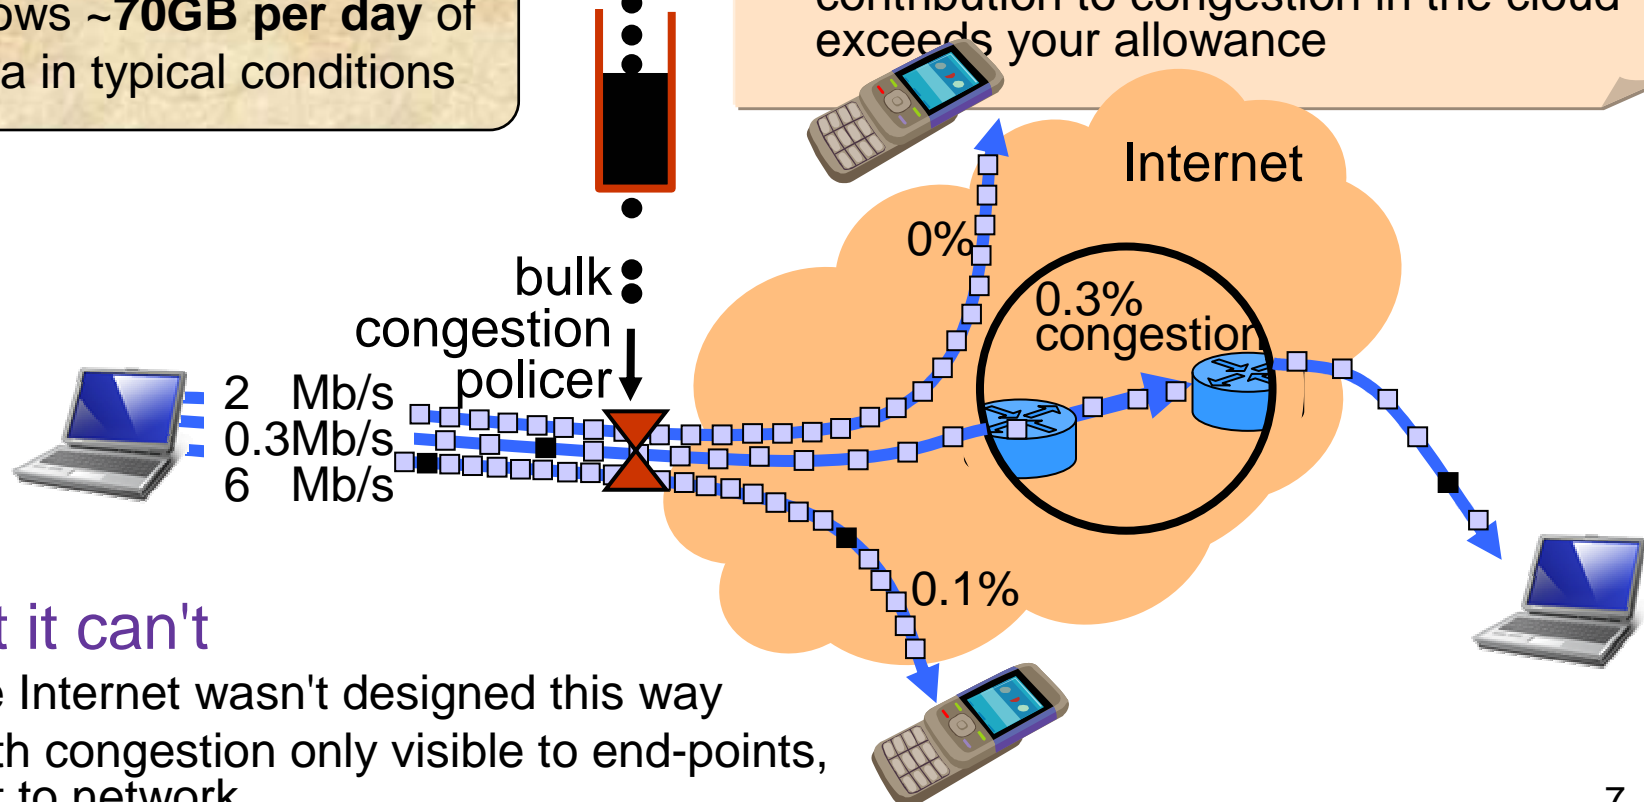
Acceptable Use Policy

'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

weighted congestion controls



- 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
- LEDBAT: a fixed example

symptoms of a lack of metric



- TCP-friendly greatly and unnecessarily restricts
 - imagine hi-speed and multipath without this restriction
- volume capping unnecessarily restricts
 - caps set to avoid even when there's no congestion to avoid

fair capacity sharing – a huge responsibility

- getting this right
 - will open a new chapter of Internet innovation
- getting it wrong
 - leaves ISPs no choice but to close off the future
 - as competition intensifies caps → app-discrimination
 - otherwise simple rate limits hurt interactive apps

more info

Re-architecting the Internet:

The Trilogy project <www.trilogy-project.org>
re-ECN & re-feedback project page:

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


These slides

<www.cs.ucl.ac.uk/staff/B.Briscoe/present.html>


bob.briscoe@bt.com

Internet resource sharing: a way forward?

discuss...



congestion volume captures (un)fairness during dynamics



main steps to deploy re-feedback / re-ECN



- 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
- requires update to the IP standard (v4 & v6)
 - started process in Autumn 2005
 - using last available bit in IPv4 header or IPv6 extension header



unilateral deployment scenarios

(non-TCP-friendly, ECN, re-ECN)


- no congestion transparency (not in protocols)
 - operator uses local congestion-volume metric in place of volume (e.g. on traffic control boxes)
 - end-host acts as if congestion-volume is limited
 - appears as voluntary as TCP, but unlikely to happen?
 - cf. BitTorrent, Microsoft & LEDBAT
- congestion transparency
 - re-ECN sender proxy

deployment scenarios

(non-TCP-friendly, ECN, re-ECN)

- academic networks and hi-speed data transfer
 - start with no policing & just conservatively weighted cc?
 - require IPv6 to have congestion policing framework?
 - sufficient proof of concept to move v4 from experimental?
 - remove of ad hoc controls when add congestion policing
- cellular networks
 - terminals & networks standardised monolithically
 - operators motivated to police heavy users [re-ECN06, re-ECN09]
 - mobile devices cross-fertilise fixed networks
 - requires radio resource control to trigger L3 ECN [Siris03]
- co-ordination
 - top-down: Global Information Infrastructure Commission (GIIC) & Internet Governance Forum (IGF)
 - as a way to distinguish net neutral behaviour from not
 - bottom-up: MIT interconnection w-g
- sticking points are bound to appear under each one

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

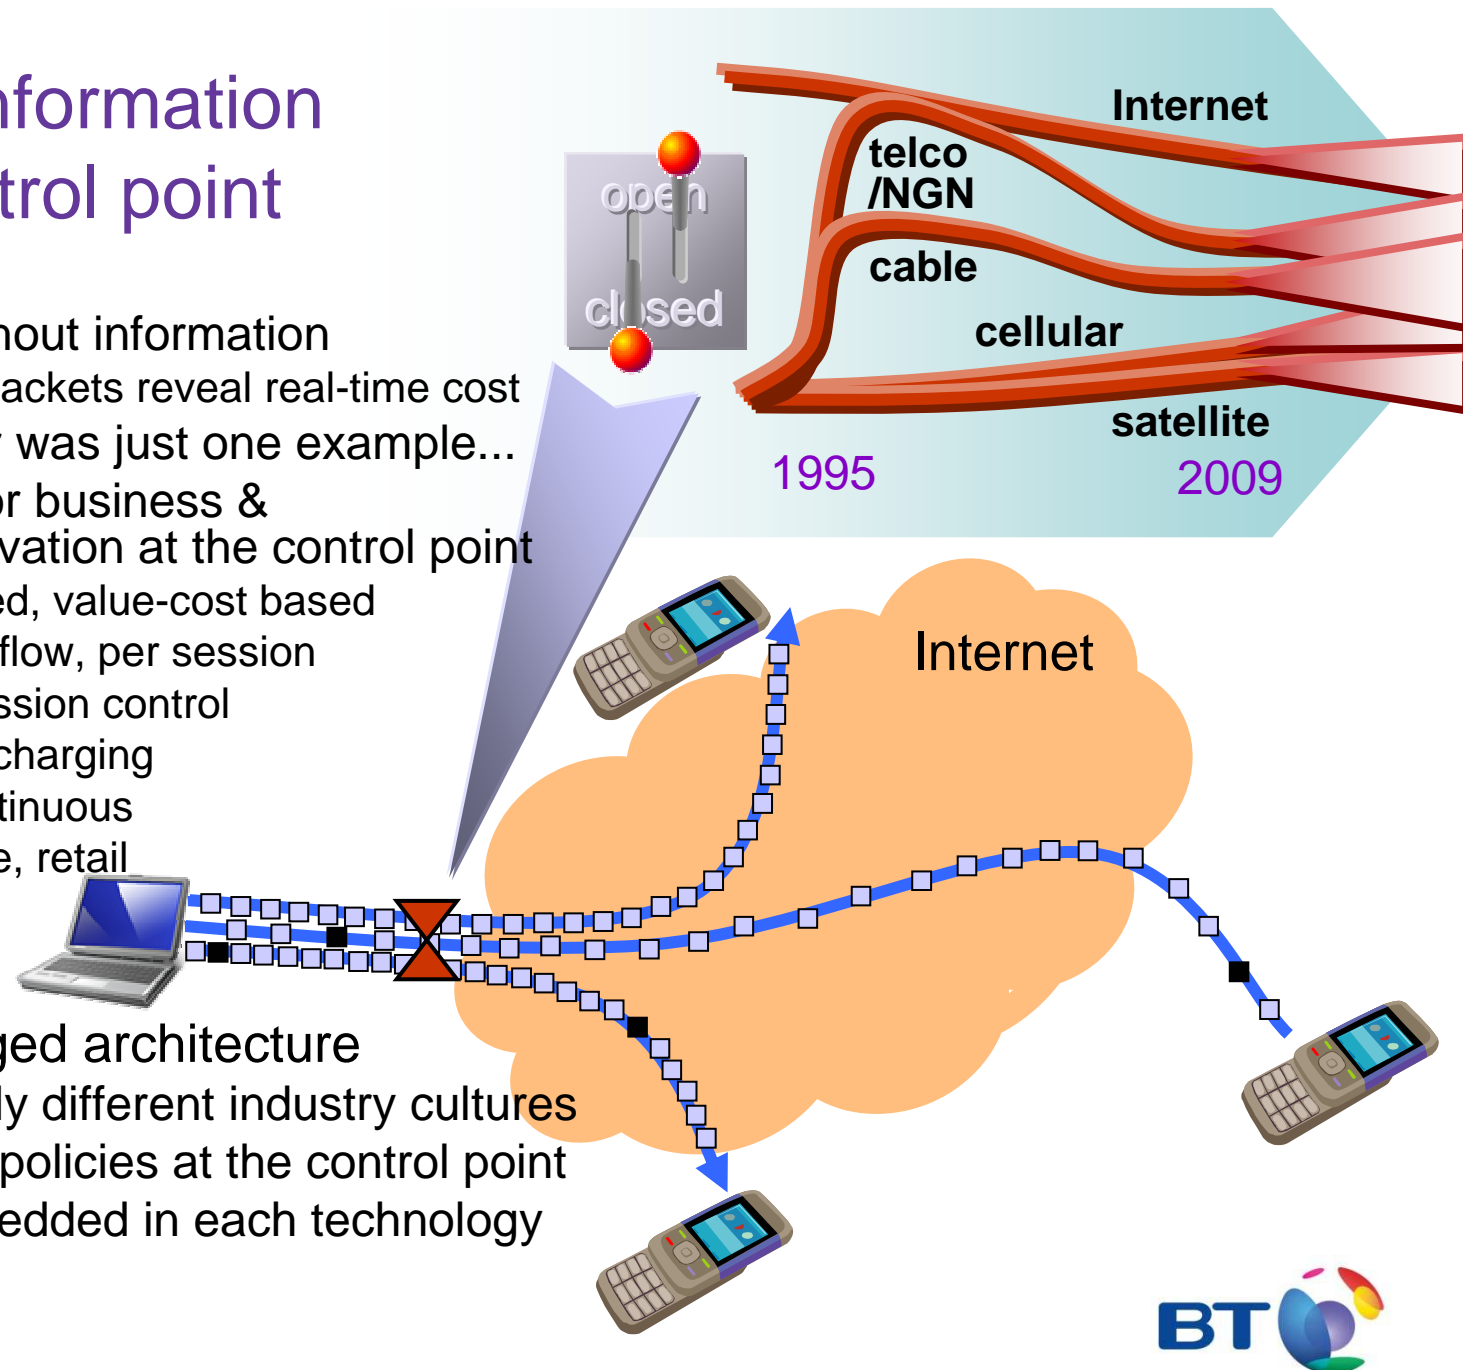


% figures =
no. of videos
that fit into the
same capacity

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

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

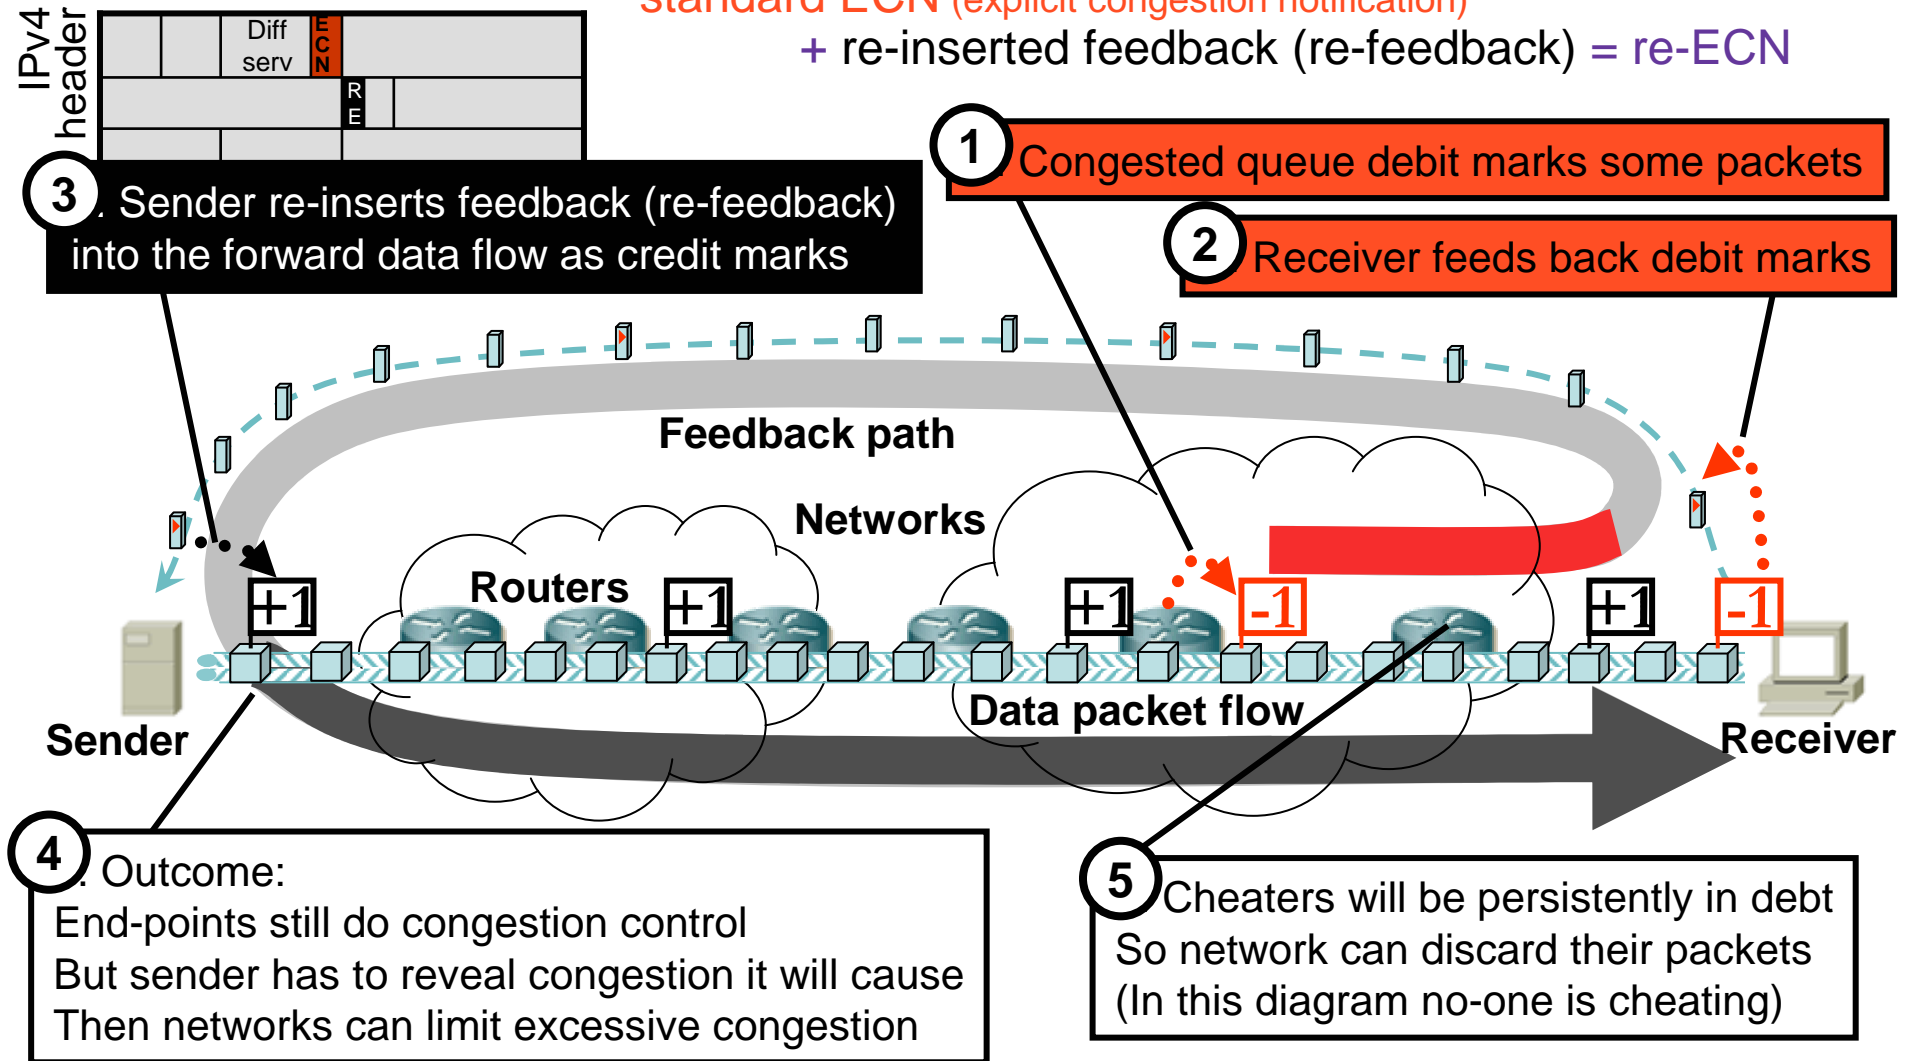


one bit opens up the future



standard ECN (explicit congestion notification)

+ re-inserted feedback (re-feedback) = re-ECN



no changes required to IP data forwarding