peer-to-peer (p2p) value & cost

Bob Briscoe
Chief Researcher
BT Group
Sep 2008

Content is King or The Long Tail?

community & social networking, interest groups

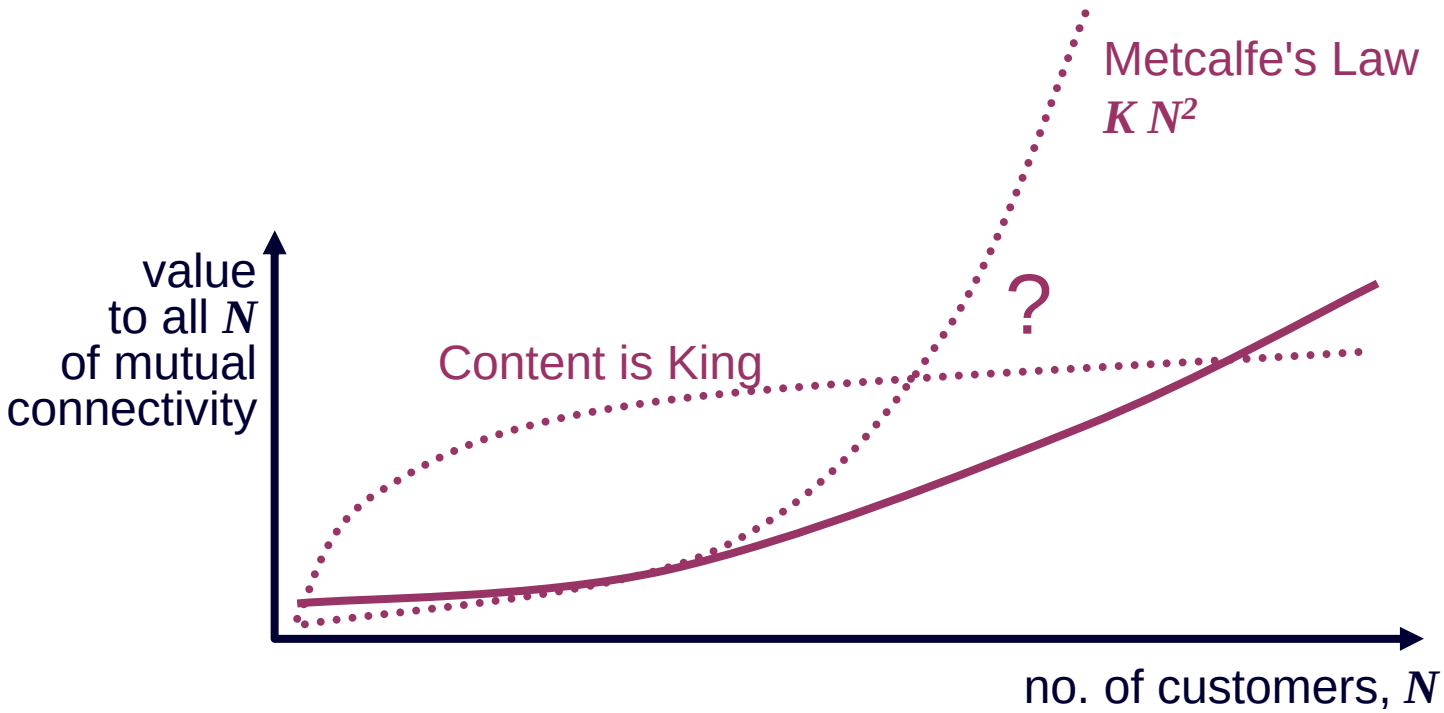

- the long tail effect eventually predominates
- but not as strongly as Metcalfe's Law predicted

Odlyzko, "Content is Not King"

Briscoe, Odlyzko & Tilly, "Metcalfe's Law is Wrong"

BT

potential peers: value in numbers

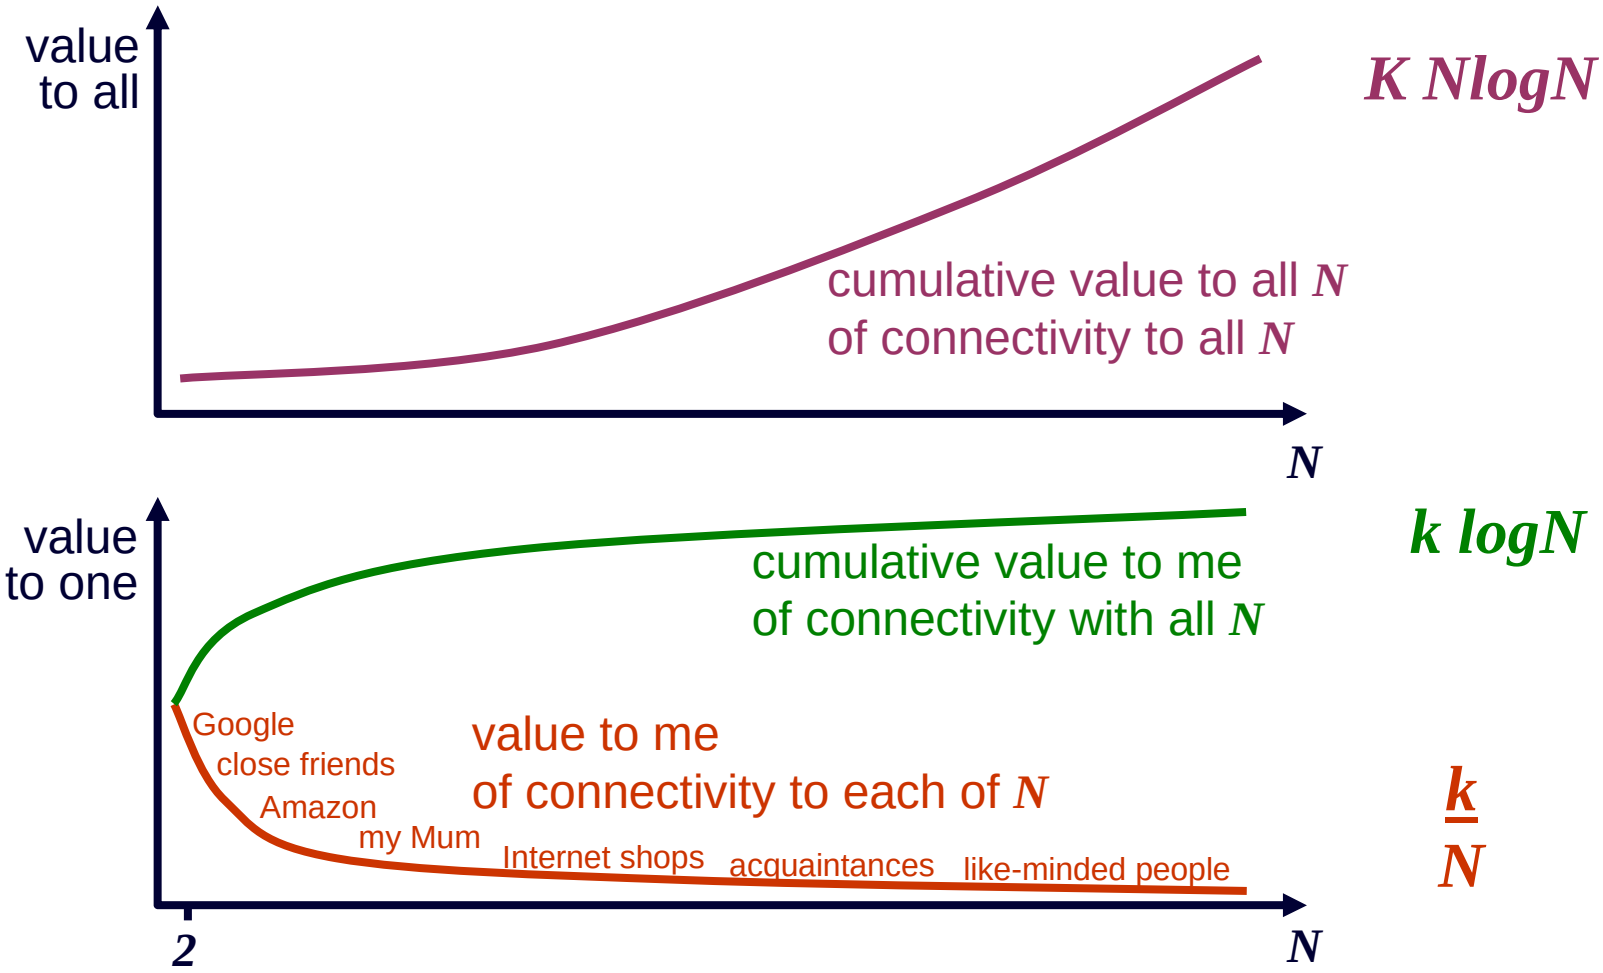

growth in potential network value

by scaling & interconnect

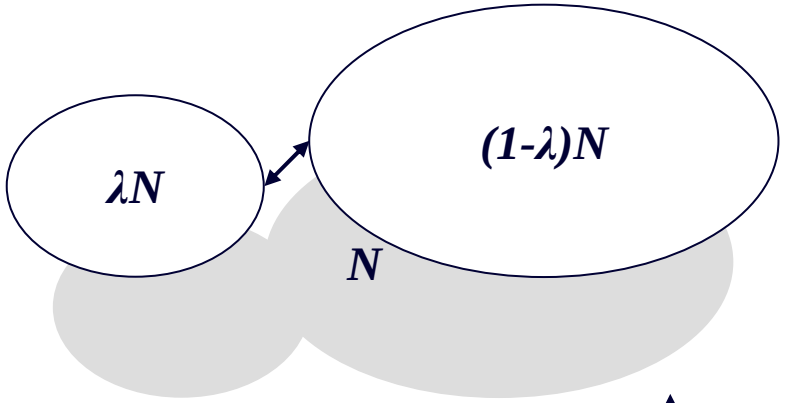

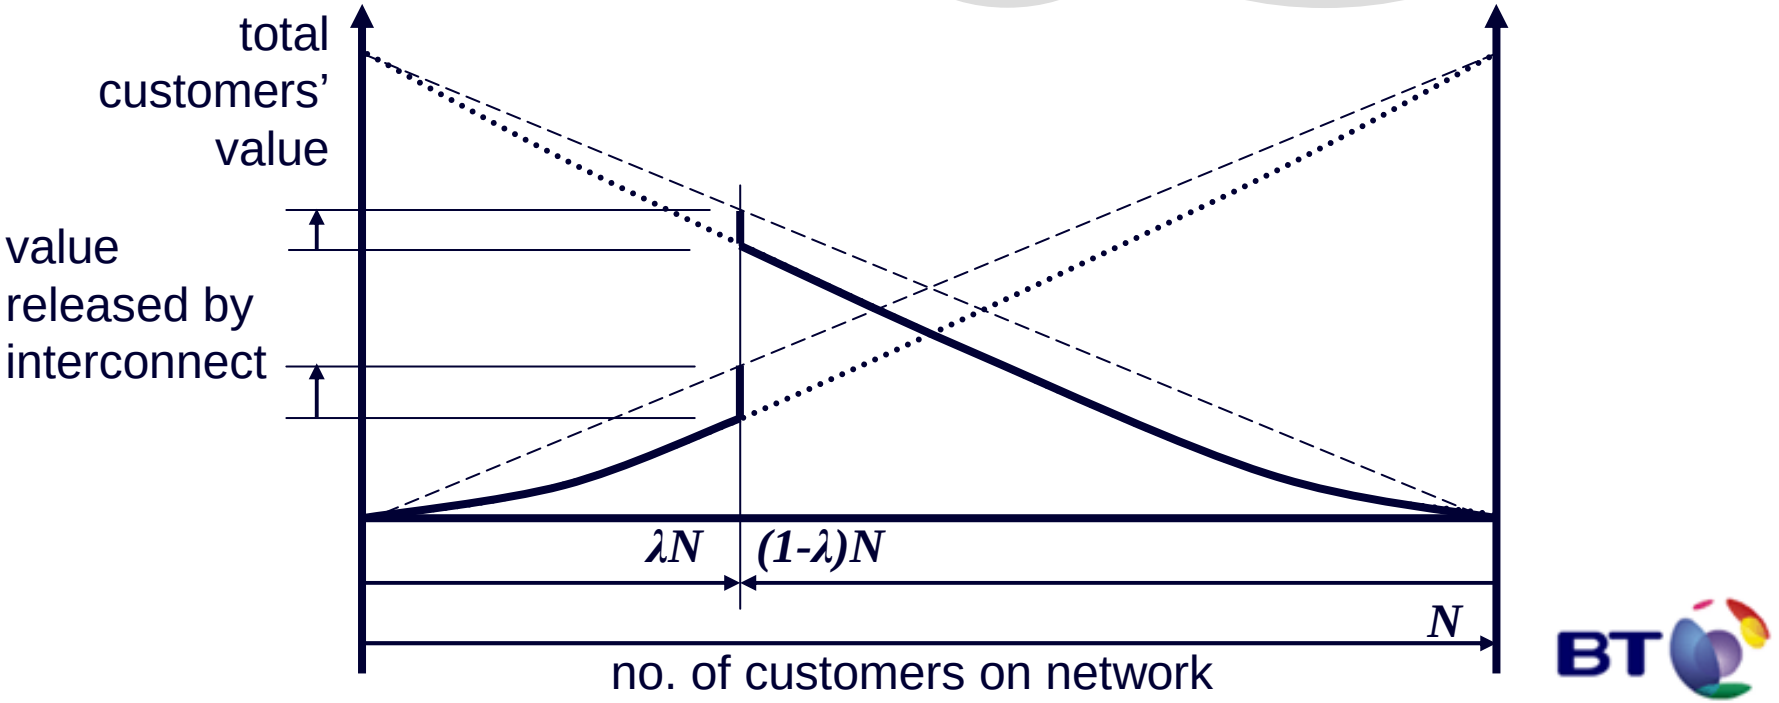

- that's all about value *to customers*
- before we start dividing the spoils between us
- remember... competition
  - drives revenue towards cost
  - ensures *customers* get the surplus value

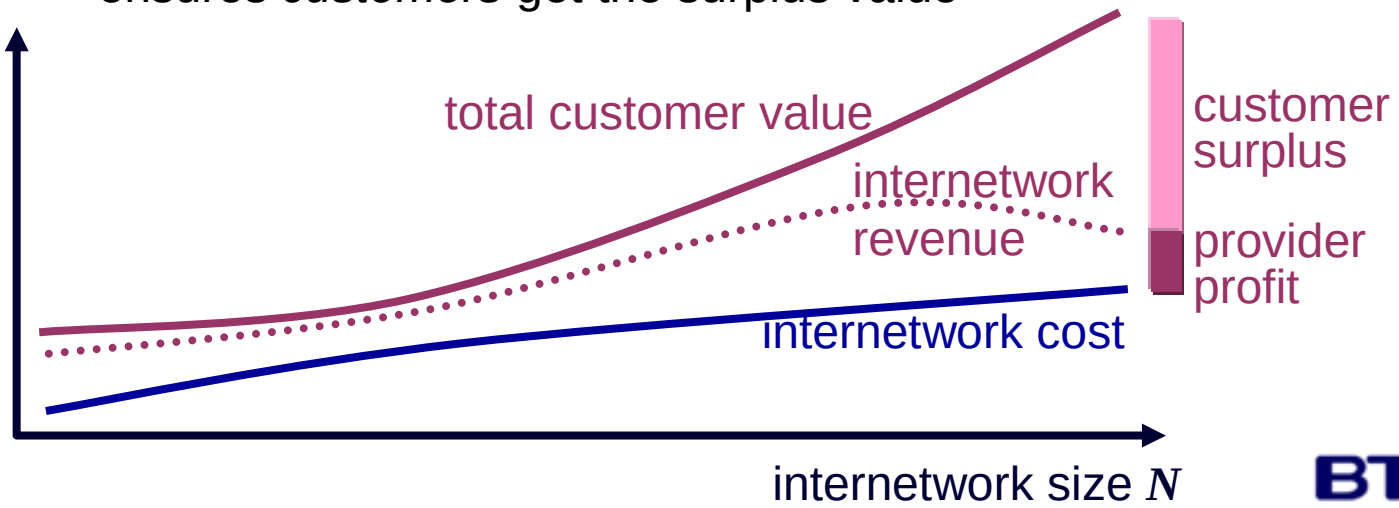

BT

the cost of p2p file-sharing?

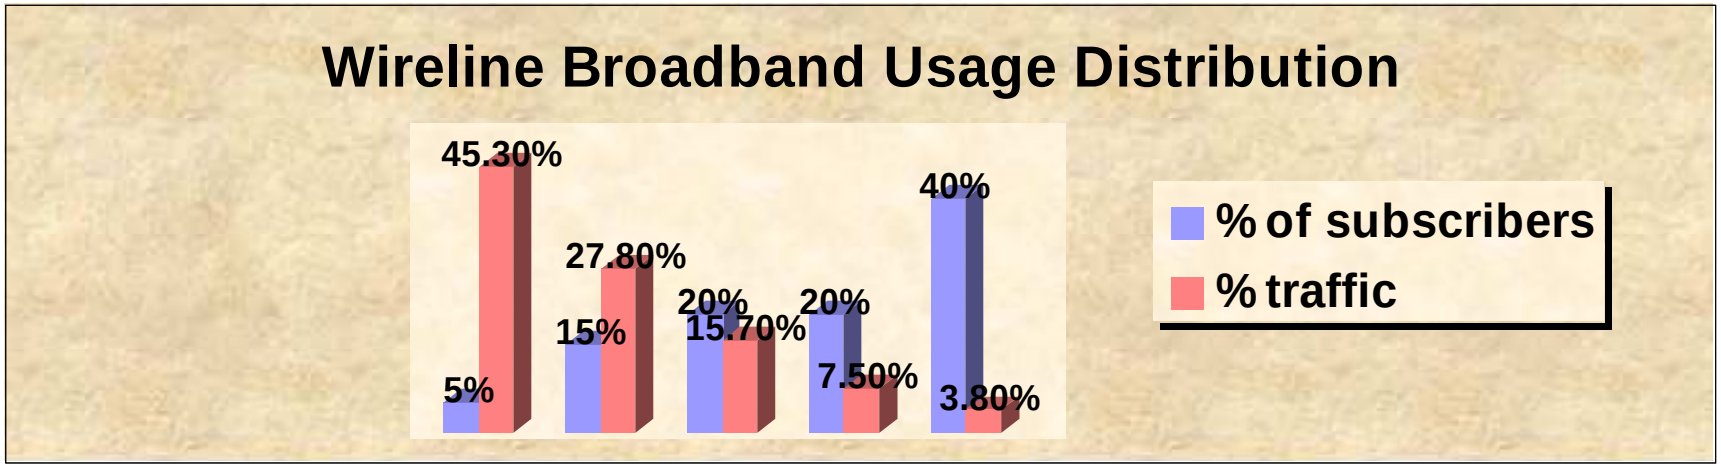

source: Ellacoya 2007
(now Arbor Networks)

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)

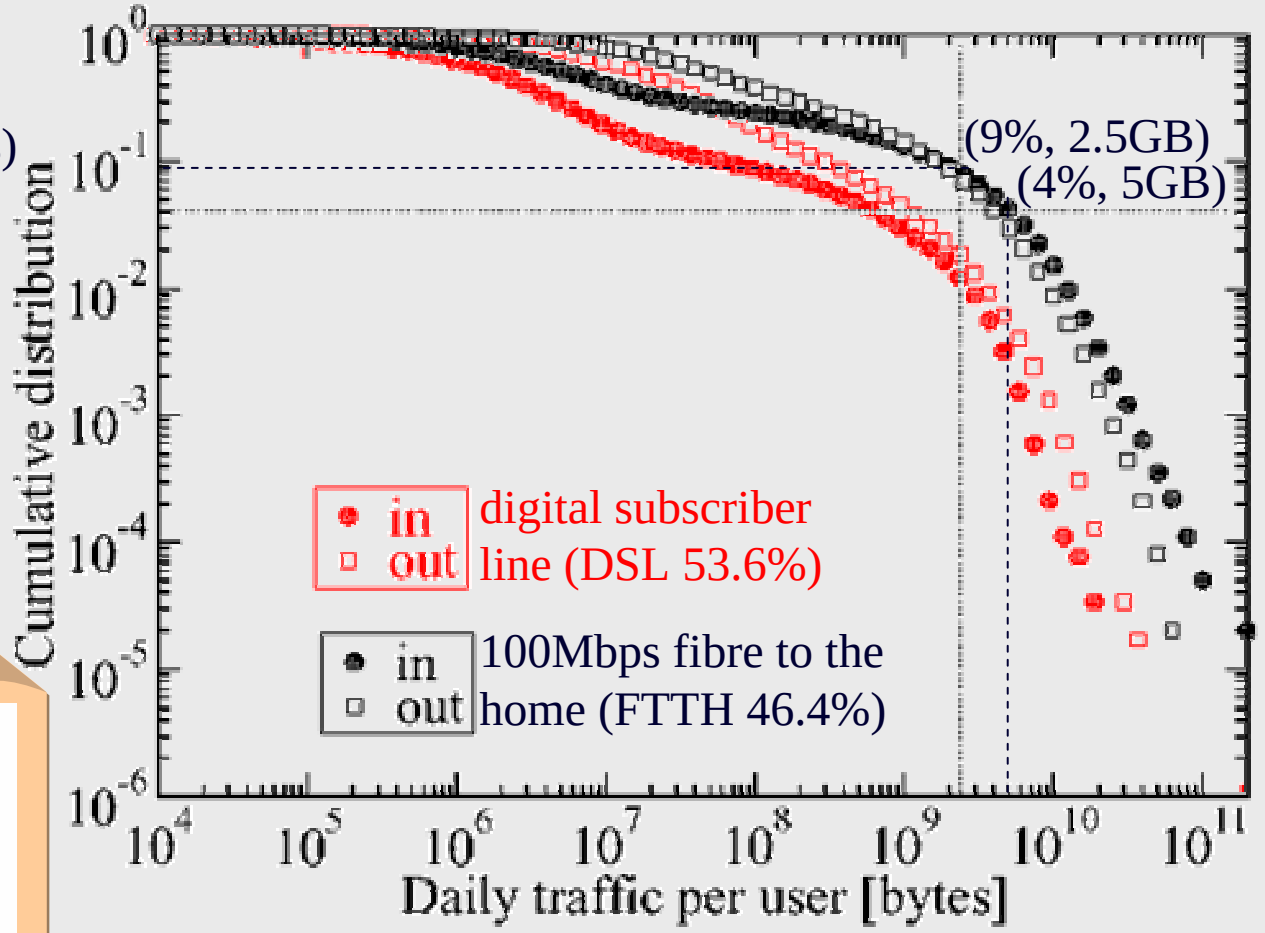

Changing technology shares of Japanese access market

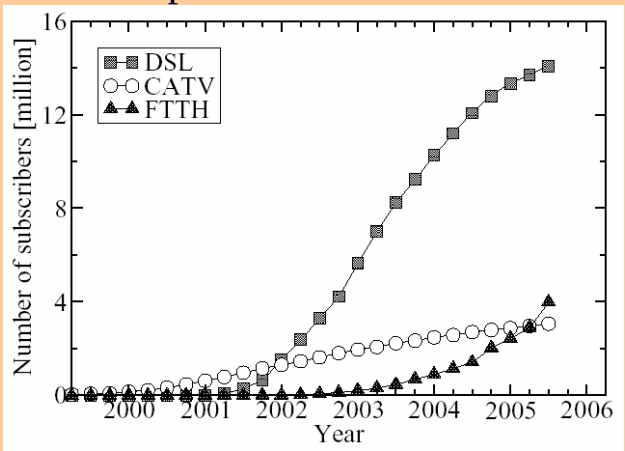

Courtesy of Kenjiro Cho et al
The Impact and Implications of the Growth in Residential User-to-User Traffic, SIGCOMM (Oct '06)

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 just takes as much best-efforts bandwidth they need
  - because of how Internet sharing 'works'
- cost of heavy usage service subsidised by ISP's lighter users
- knee-jerk reaction of ISP
  - block p2p or independent services
- No! don't blame your customers
- fix the cost accountability foundations
  - separation between network & services is good
  - but need to add cost accountability to IP

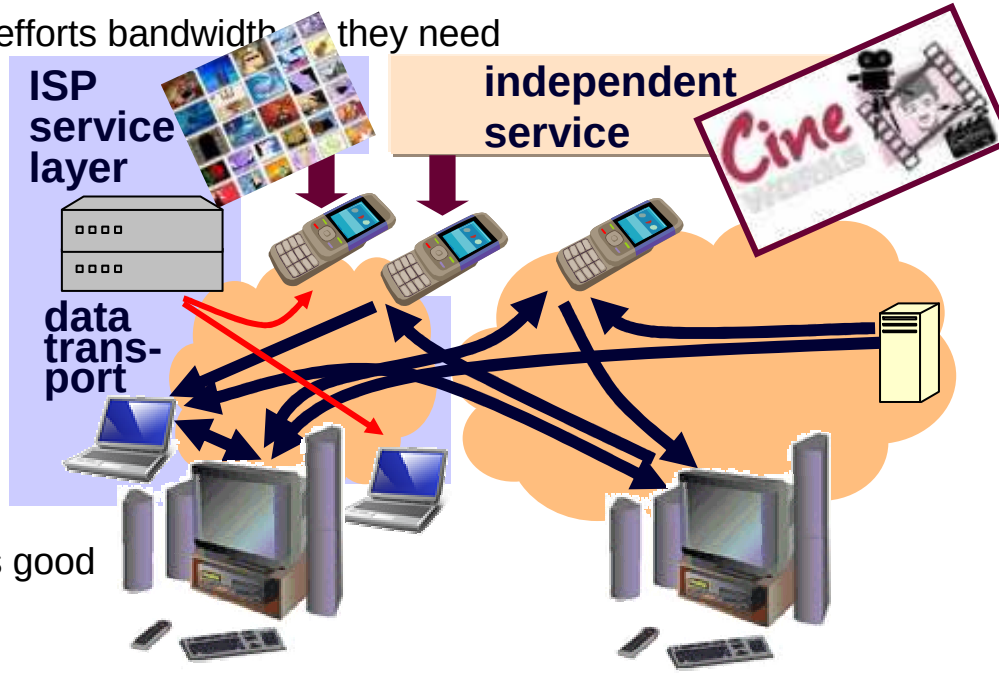

BT

underlying problems
blame our choices, not p2p

- commercial
  - Q. what is cost of network usage?
  - A. volume? NO; rate? NO
  - A. 'congestion volume'
  - *our own* unforgivable sloppiness over what our costs are
- technical
  - lack of cost accountability in the Internet protocol (IP)
  - p2p file-sharers exploiting loopholes in technology *we've* chosen
- *we* haven't designed *our* contracts & technology for machine-powered customers

BT

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

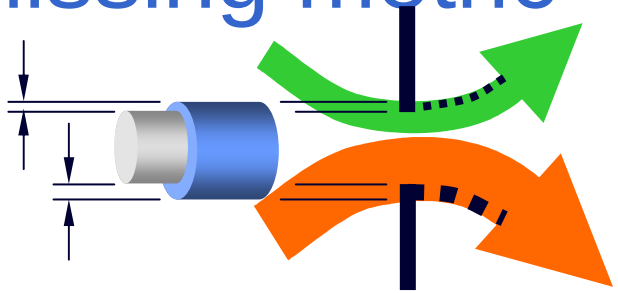

note: diagram is conceptual

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

BT

how Internet sharing 'works'

endemic congestion & voluntary restraint

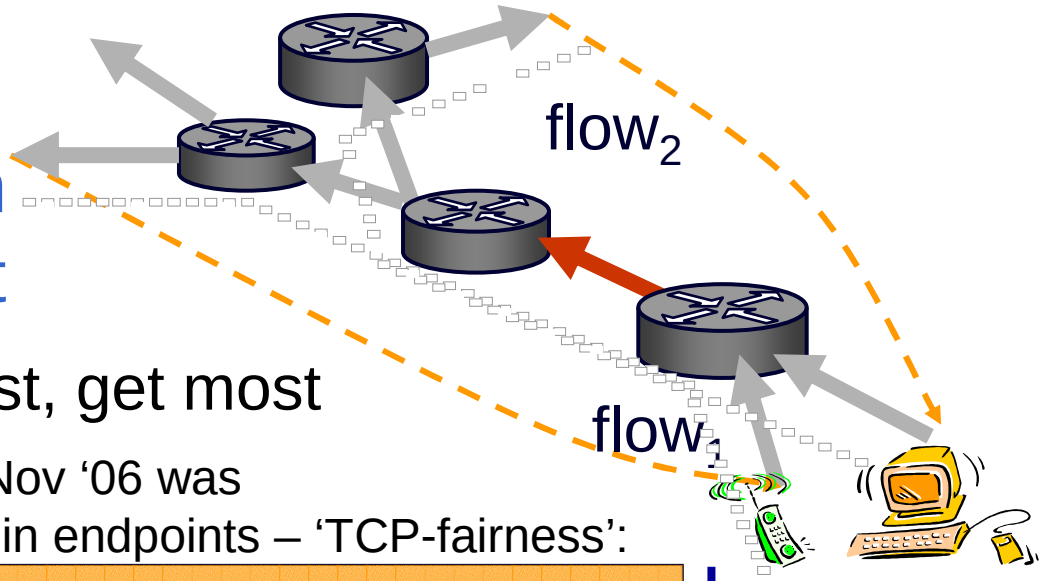

- aka. those who take most, get most
  - technical consensus until Nov '06 was voluntarily polite algorithm in endpoints – 'TCP-fairness':

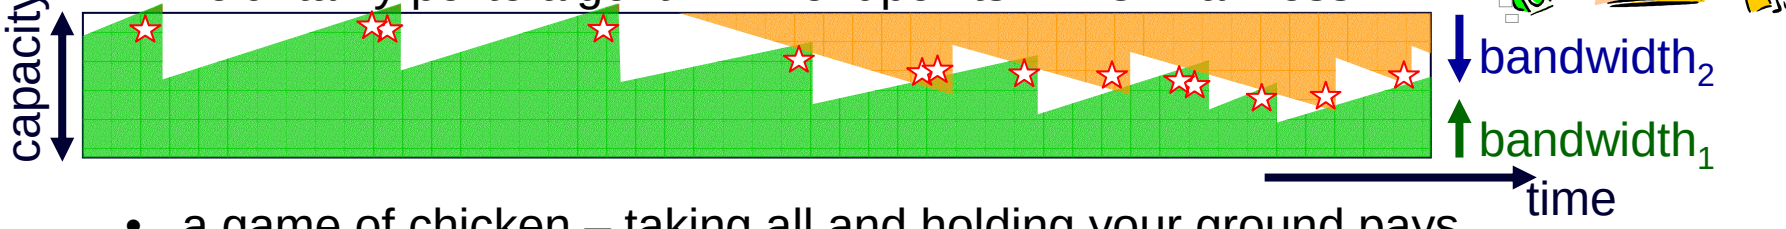

  - a game of chicken – taking all and holding your ground pays

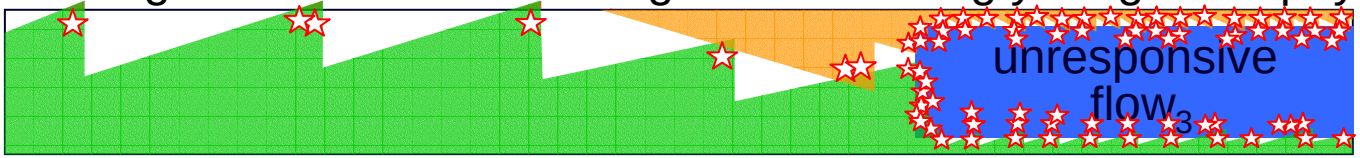

(VoIP, VoD Joost 700kbps)

  - 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)

BT

fairer is faster

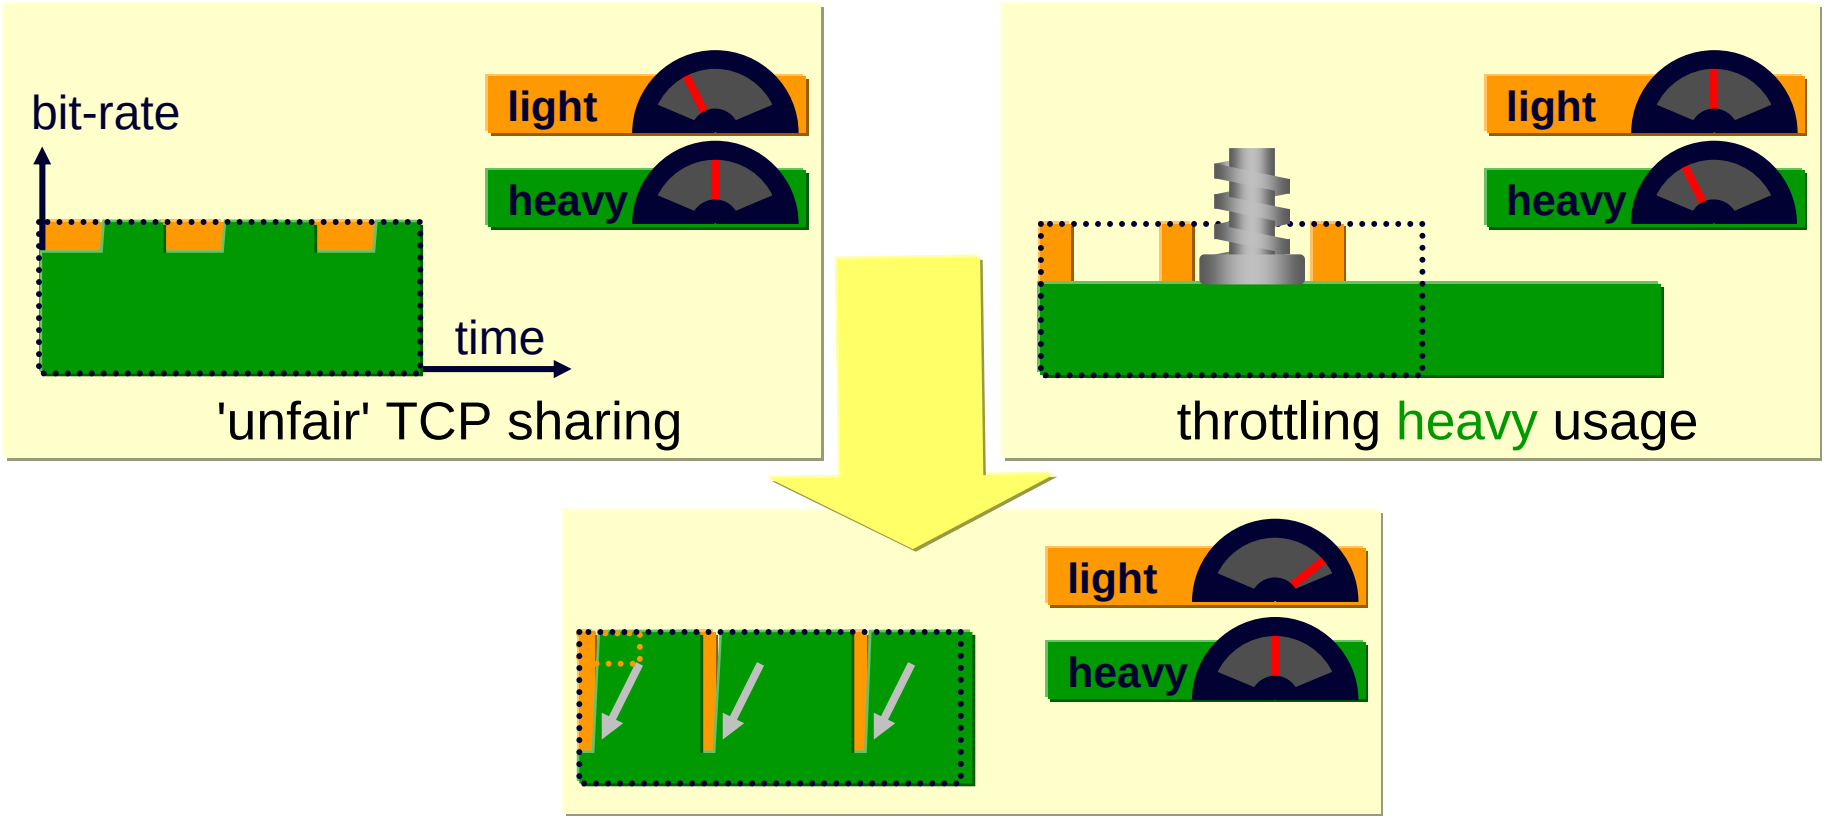

- what's required: limit congestion, not volume
  - then heavy usage will back away whenever light usage appears
  - so light usage can go much faster
  - hardly affecting completion times of heavy usage

BT

limiting congestion?

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

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

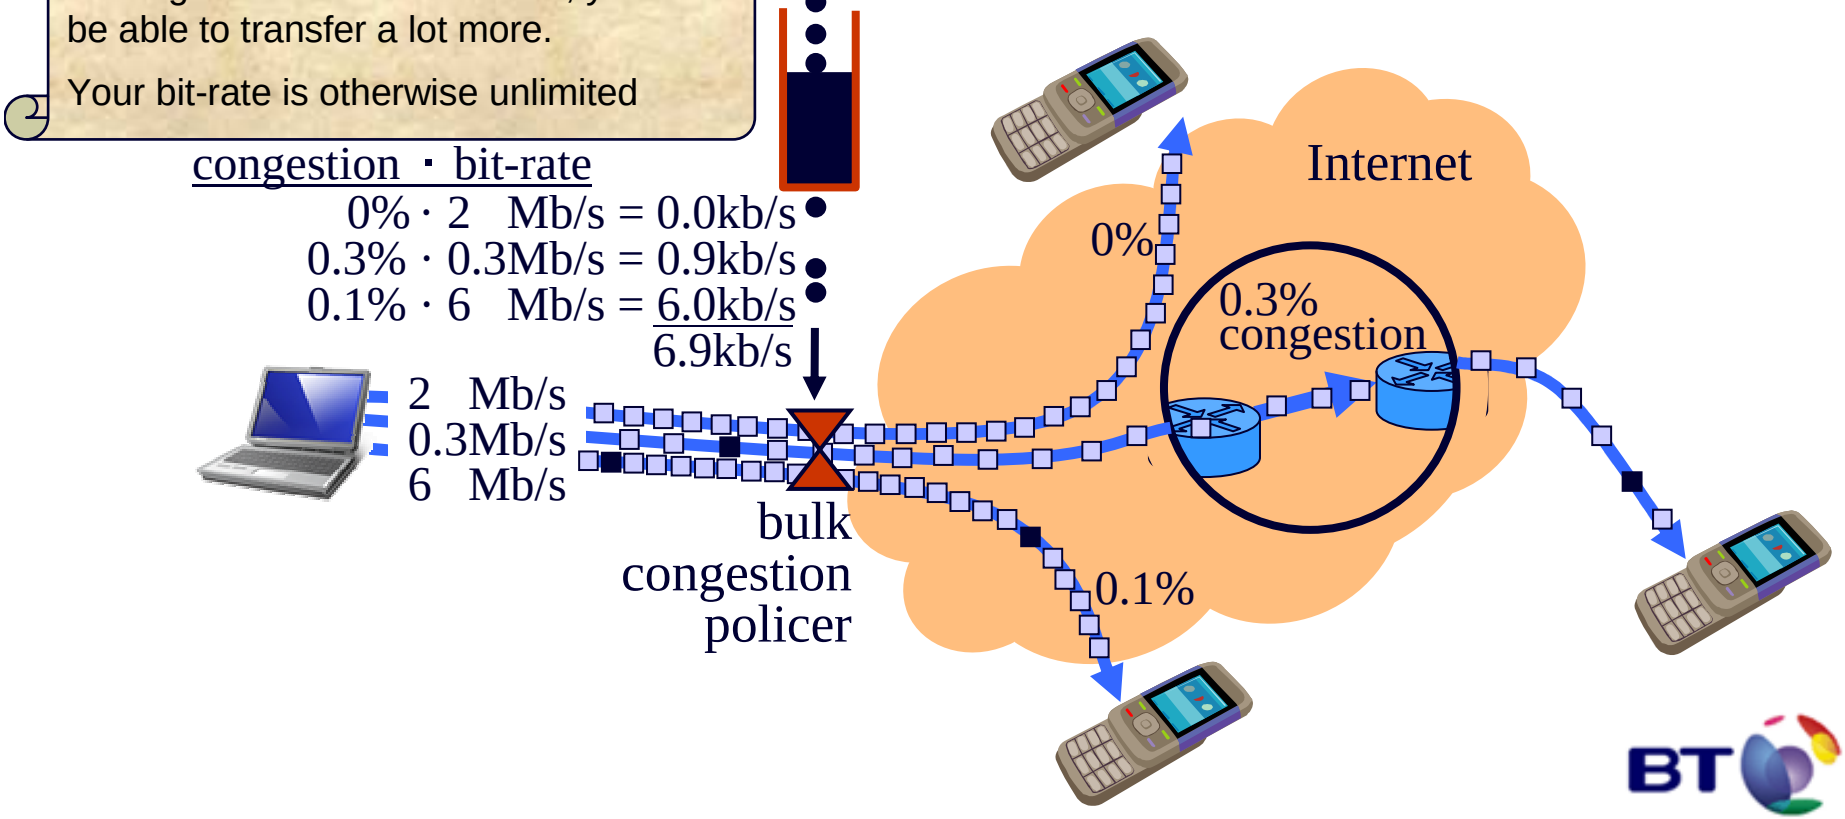

problems using congestion in contracts

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

BT

automatic
usage cost allocation

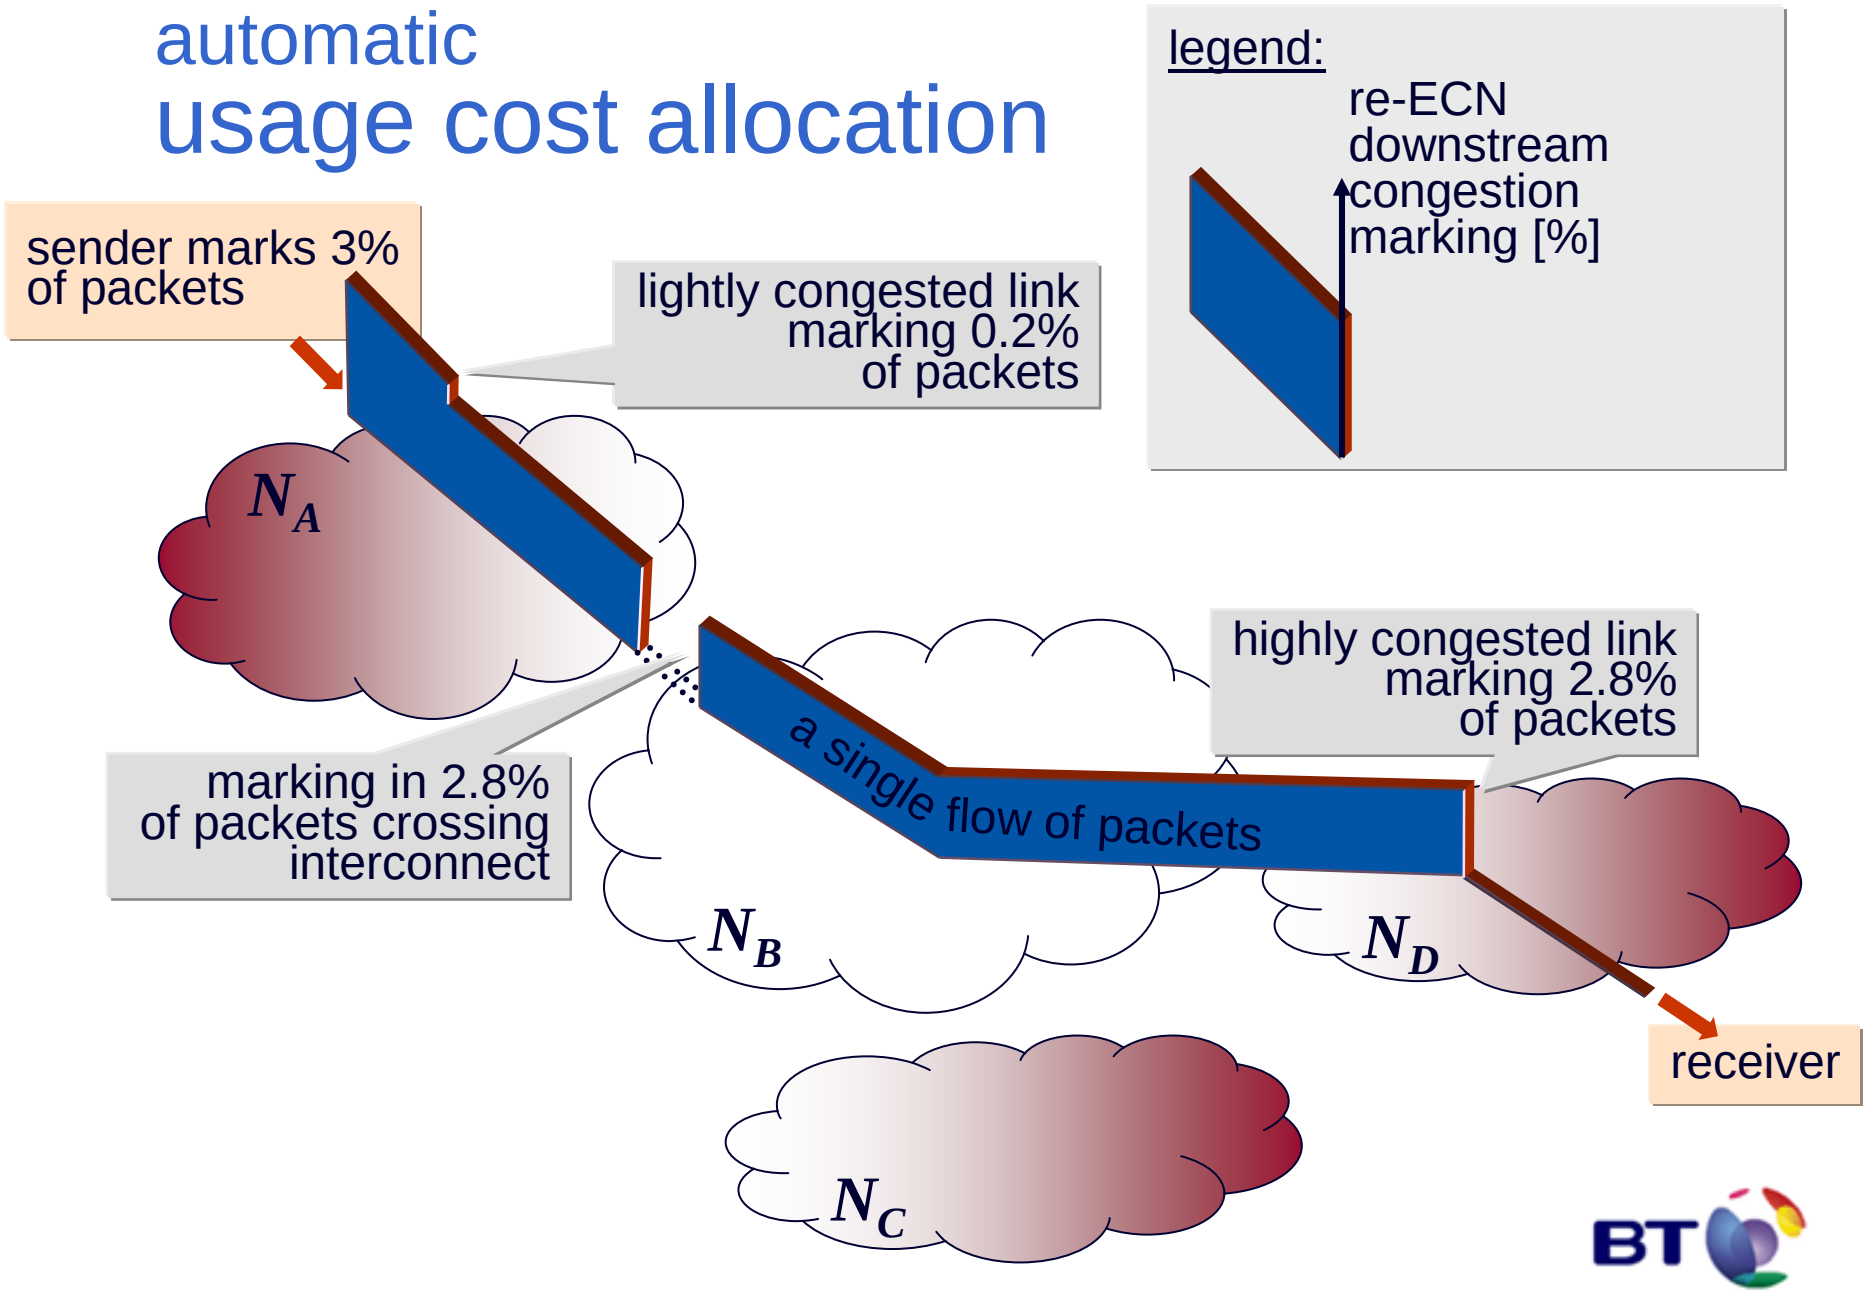

interconnect aggregation

simple internalisation of all externalities

'routing money'

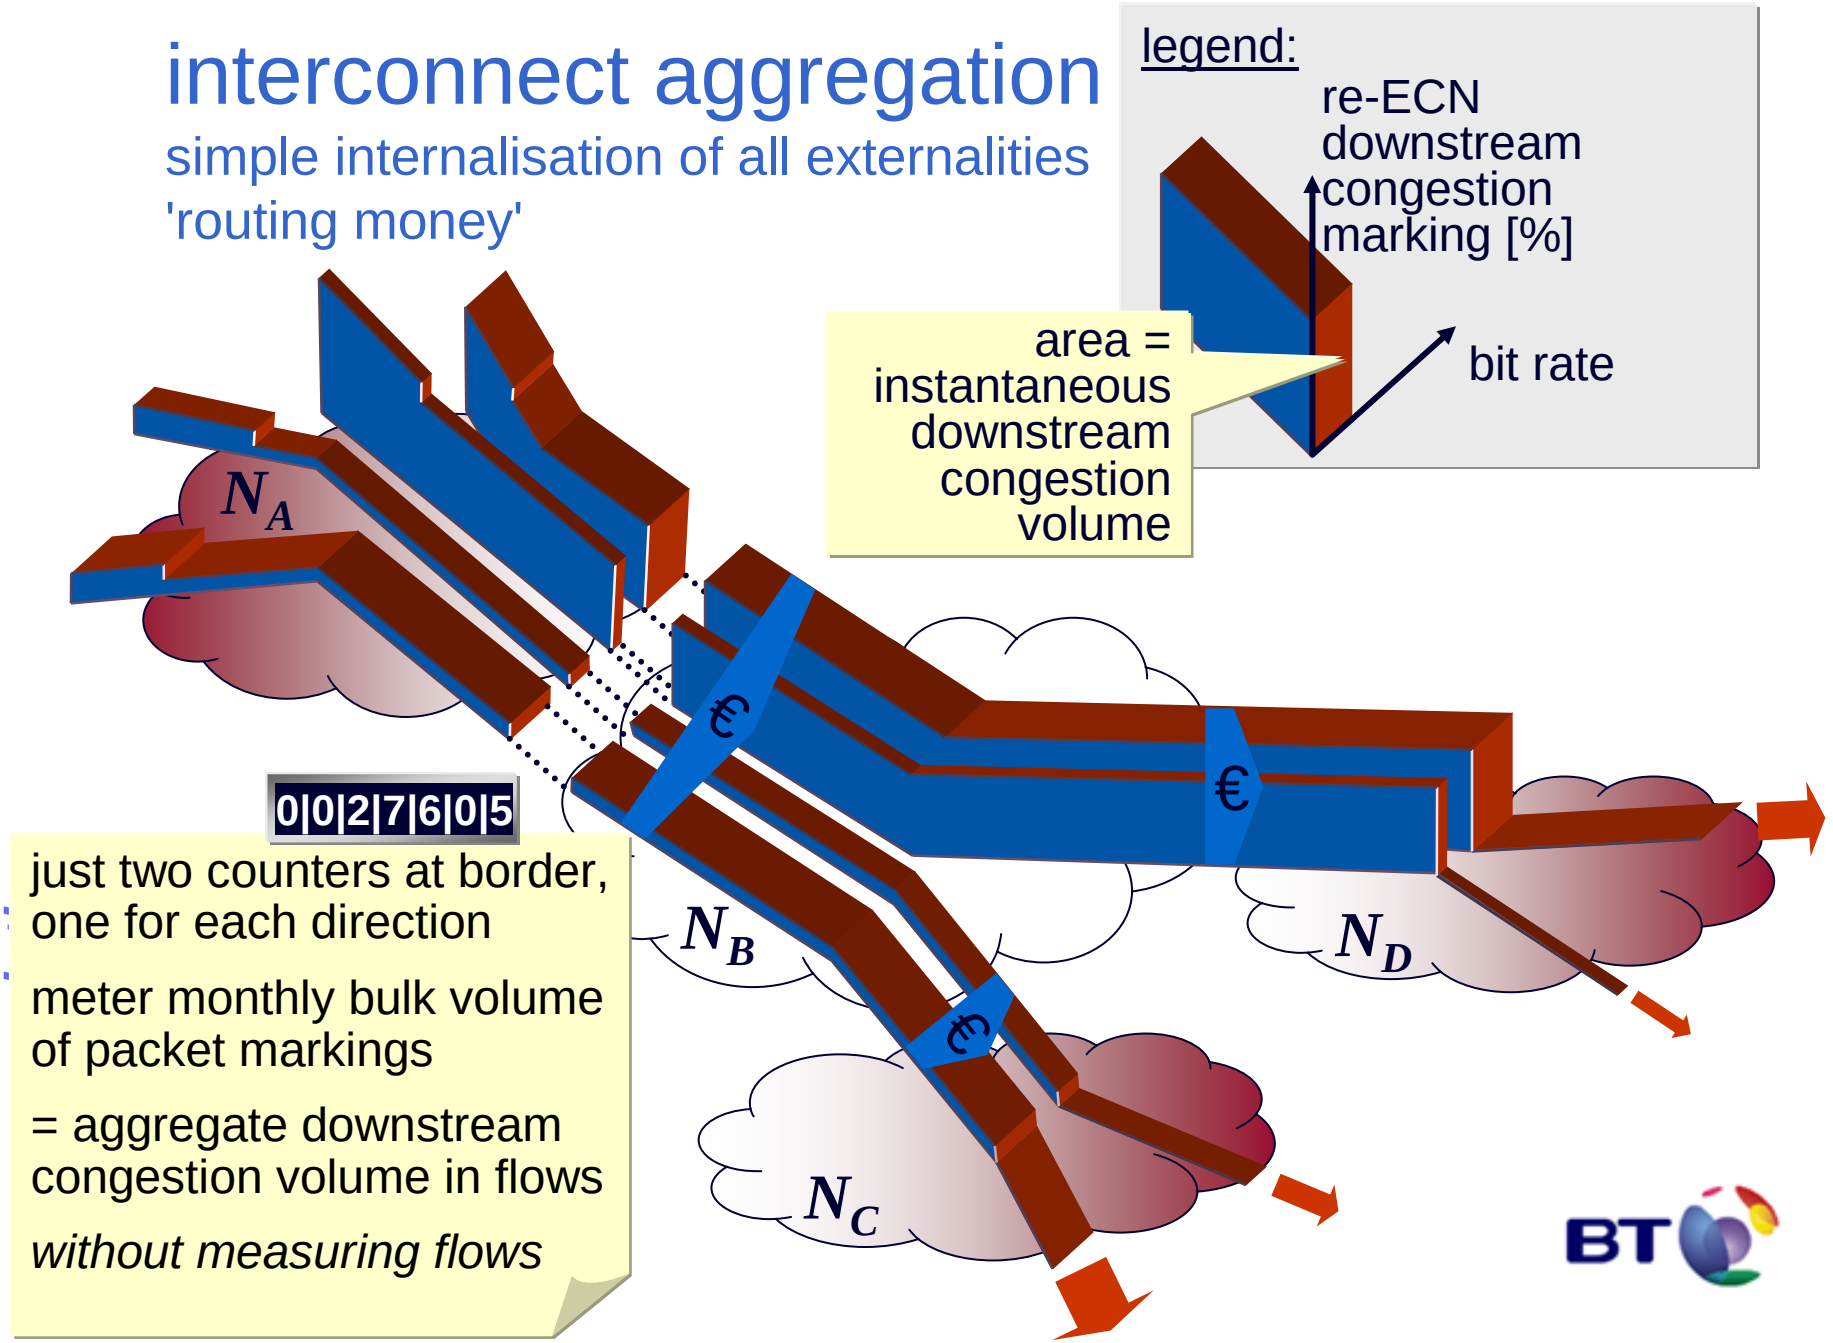

just two counters at border, one for each direction

meter monthly bulk volume of packet markings

= aggregate downstream congestion volume in flows

without measuring flows

example sustainable business model

for basic data transport

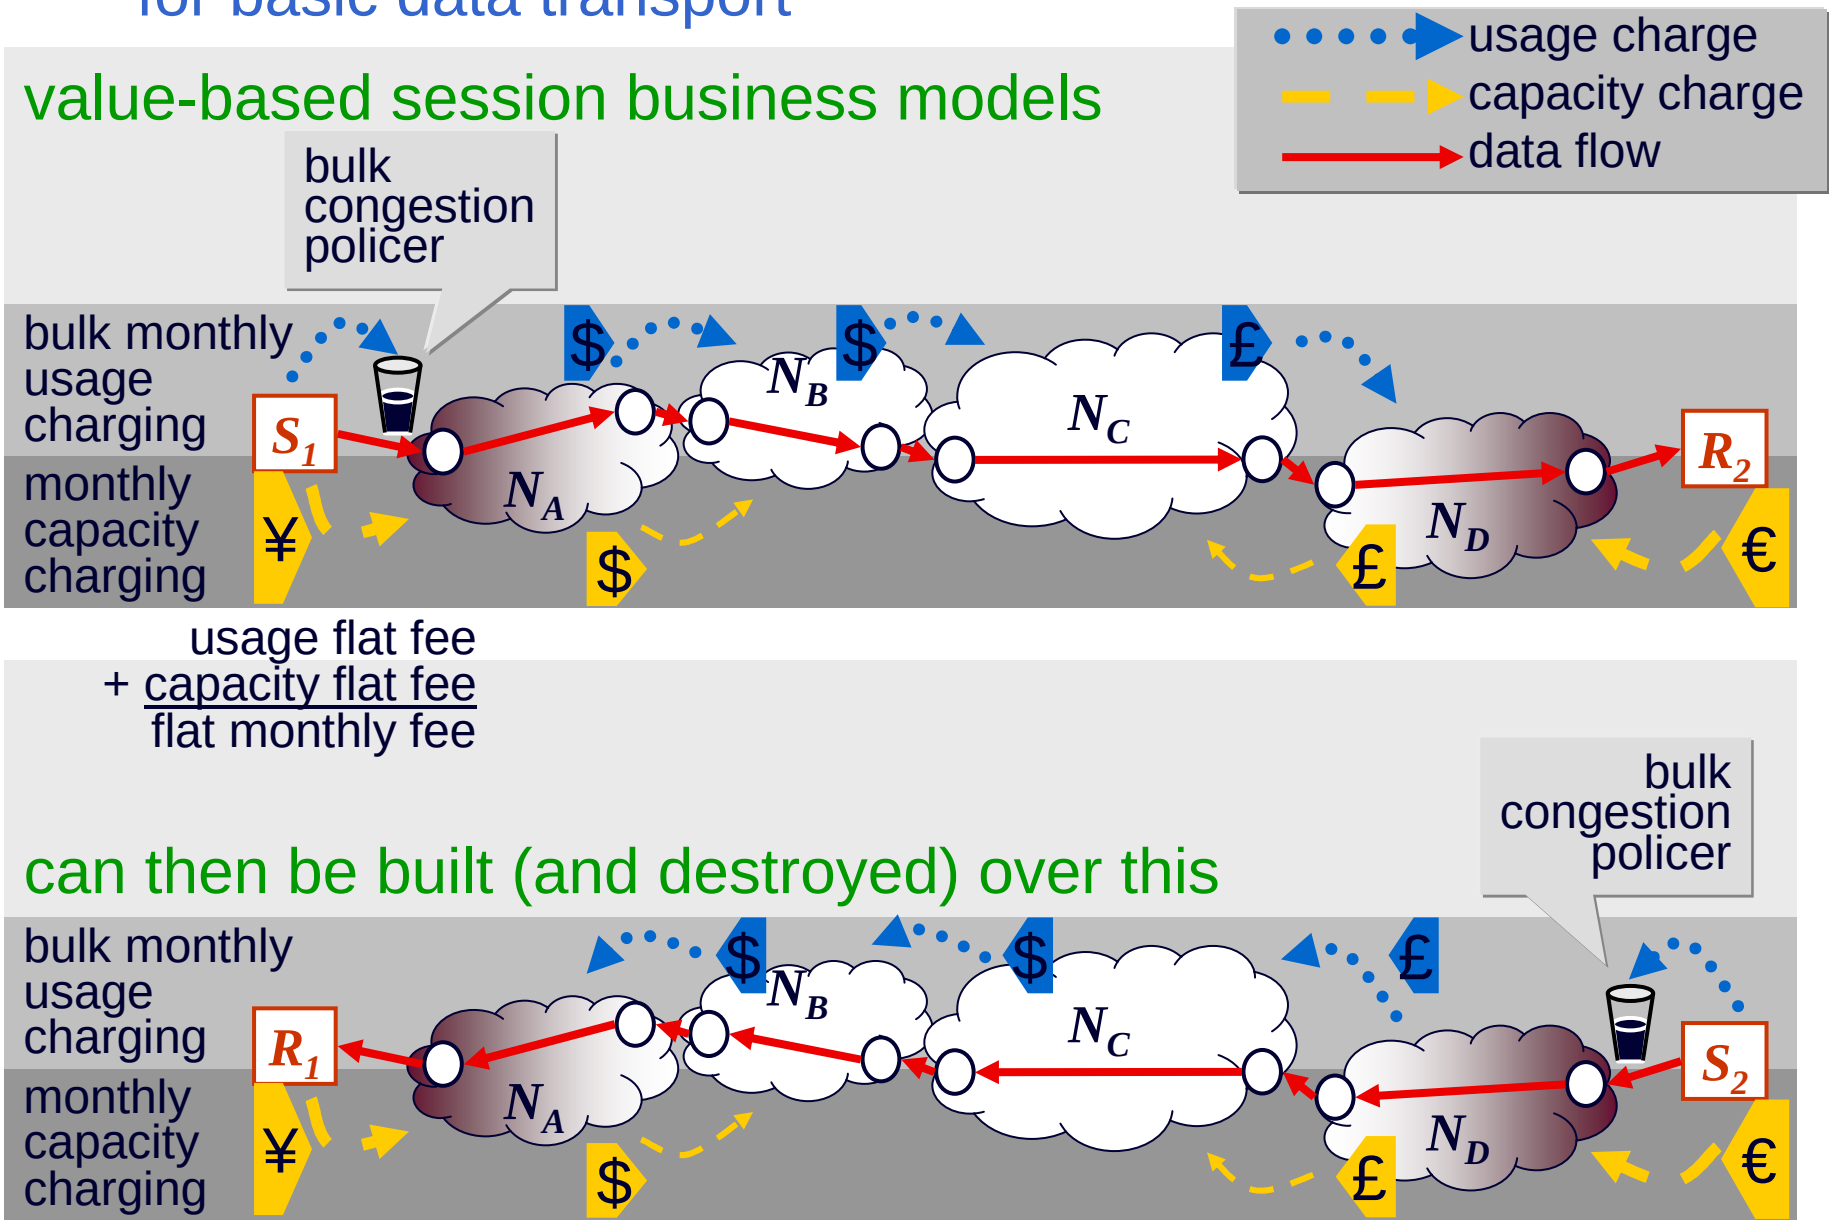

wrap up

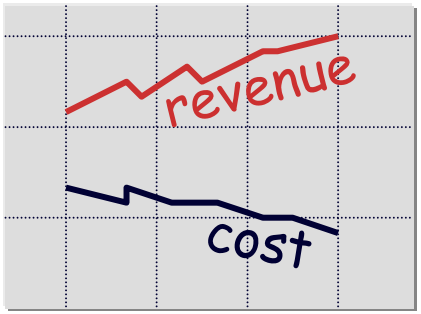

- expect consistent value growth from p2p
  - but don't just focus on value, cost as well
- separation of service & network is excellent industry goal
  - but how TCP/IP shares cost of transport needs serious attention
  - understanding of network economics is young, so is IP
- not the fault of ISP's customers or of p2p
  - ISPs will need a mitigating strategy until it's fixed
- please help us add cost accountability (re-ECN) to IP
  - please brief your technical & standards strategy people
- a platform on which customer contracts can be built
  - for basic transport, then services on top
  - so billions of machines work together in everyone's best interests

BT

more info...

- Growth in value of a network with size
  - Bob Briscoe, Andrew Odlyzko & Ben Tilly, "Metcalfe's Law is Wrong", IEEE Spectrum, Jul 2006
- Inevitability of policing
  - 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
  - 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
  - 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)
- Re-architecting the Future Internet:
  - The Trilogy project
- Re-ECN & re-feedback project page:
  <http://www.cs.ucl.ac.uk/staff/B.Briscoe/projects/refb/>

p2p
value & cost

Q&A

congestion competition – inter-domain routing

- if congestion → profit for a network, why not fake it?
  - upstream networks will route round more highly congested paths
  - N_A can see relative costs of paths to R_1 thru N_B & N_C
- the issue of monopoly paths
  - incentivise new provision
  - as long as competitive physical layer (access regulation), no problem in network layer

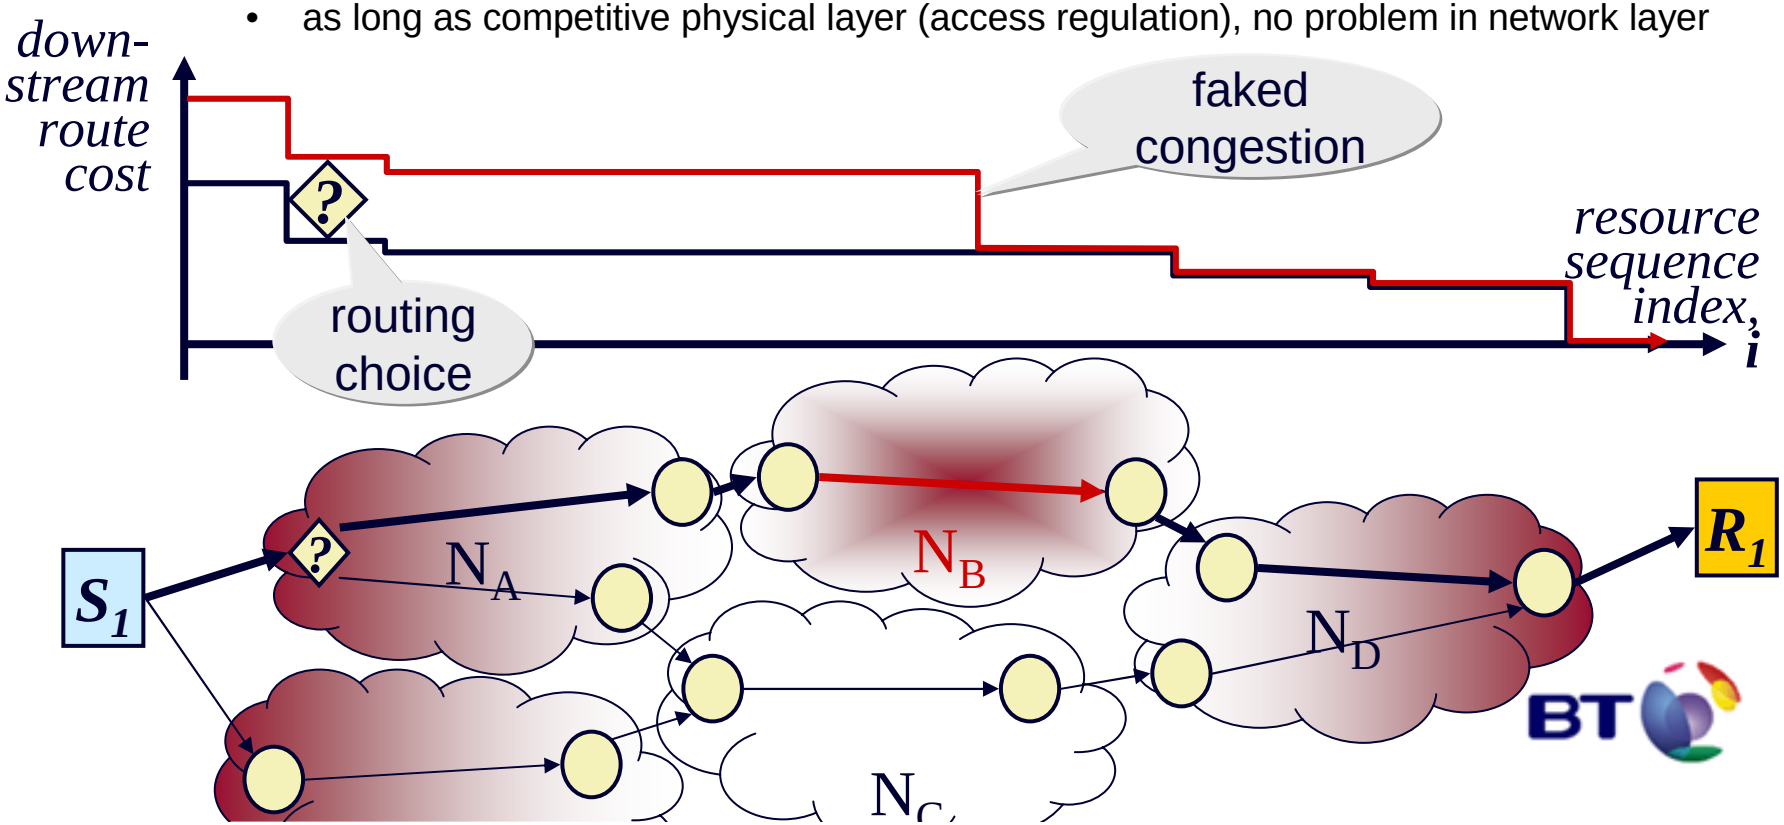

main steps to deploy re-feedback / re-ECN

- network
  - turn on explicit congestion notification in routers (already available)
  - 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
- customer contracts
  - include congestion cap
- oh, and first we have to update the IP standard
  - started process in Autumn 2005
  - using last available bit in the IPv4 packet header
  - IETF recognises it has no process to change its own architecture
  - Apr'07: IETF supporting re-ECN with (unofficial) mailing list & co-located meetings

BT