Collective delusions behind how capacity gets shared

Bob Briscoe
with Toby Moncaster & Lou Burness
presented by Dirk Trossen
Jan 2008

BT

freedom
to limit the freedom of others?

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

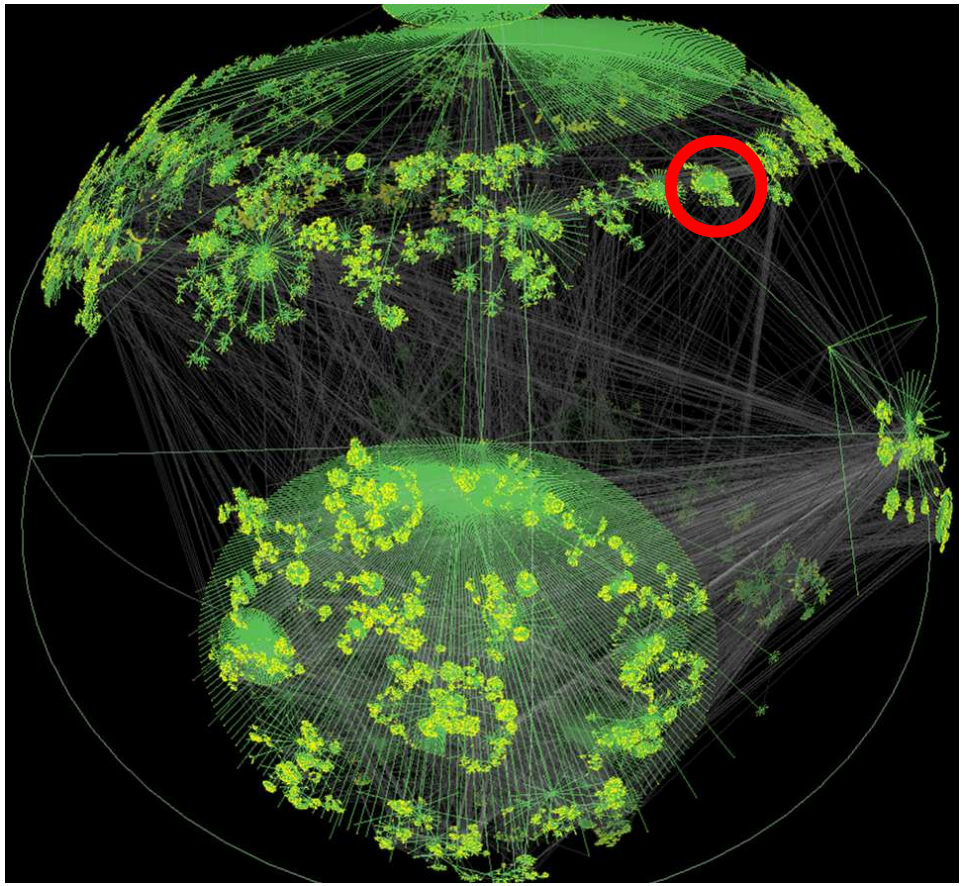

Internet topology visualization produced by Walrus (Courtesy of Young Hyun, CAIDA)

for scale: ~10M lines ringed in red

BT

fair bottleneck bit-rate?

two incompatible partial worldviews

the Internet way (TCP)	operators (& users)
‘flow rate equality’	‘volume accounting’
per data flow	per user
instantaneous	over time

this talk

- status report on our attempts to unveil multiple delusions
- the standards and research community’s double delusion
  - TCP’s equal flow rates are no longer fair at all (by any definition)
  - TCP protocol increasingly doesn’t determine capacity shares anyway

BT

base example

different activity factors

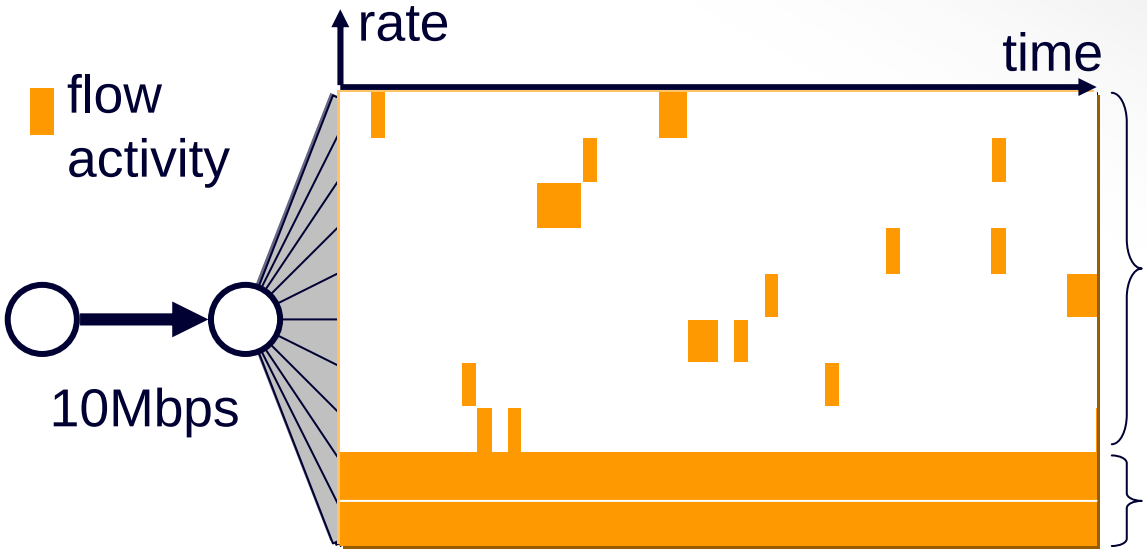

2Mbps access each

80 users of attended apps

20 users of unattended apps

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

compounding activity factor & multiple flows

no-one is saying more volume is unfair

but volume accounting says it's fairer if heavier users get less rate during peak period

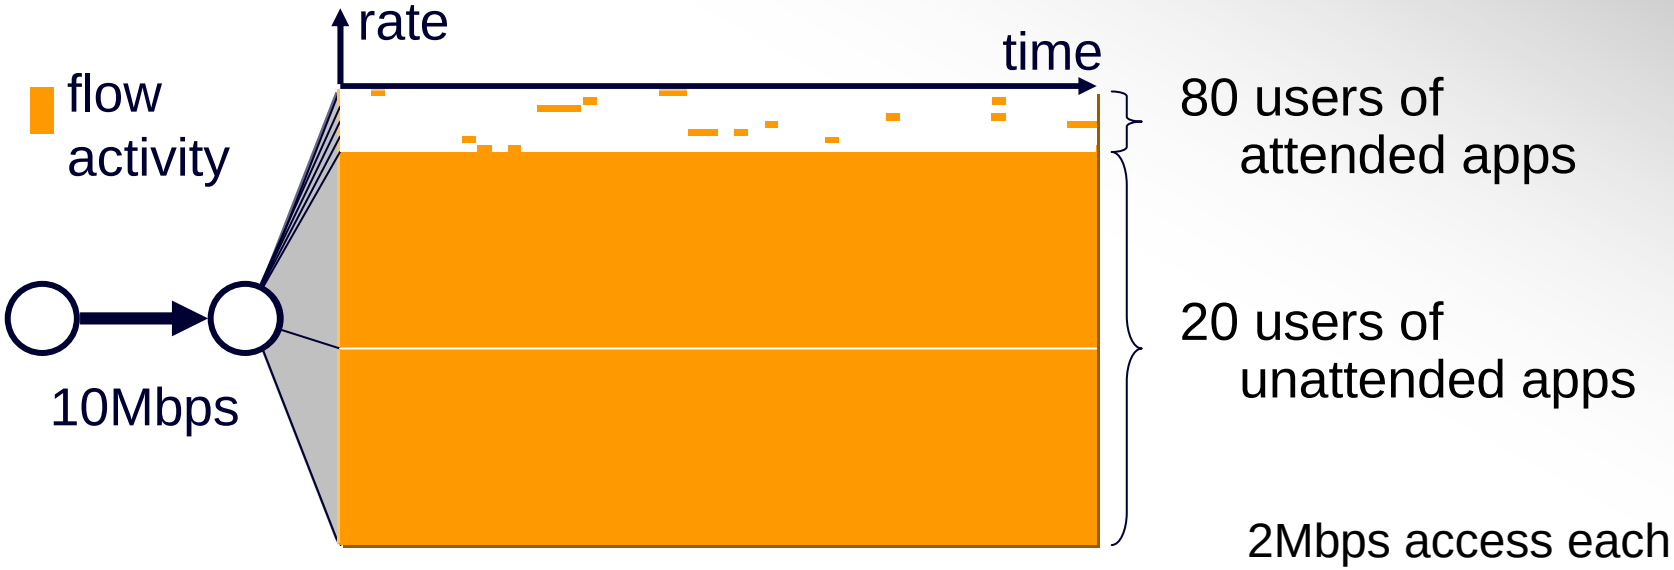

usage type	no. of users	activity factor	ave.simul flows /user	TCP bit rate /user	vol/day (16hr) /user	traffic intensity /user
attended	80	5%	2	20kbps	7.1MB	1kbps
unattended	20	100%	50	500kbps	3.6GB	500kbps
				x25	x500	x500

realistic numbers?
there are elephants in the room

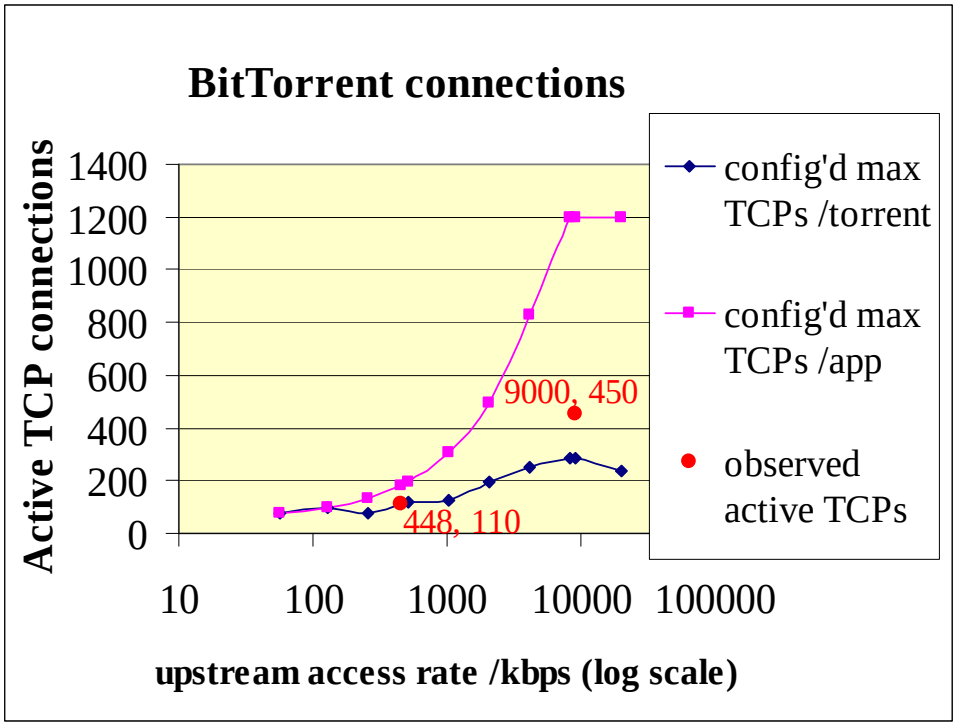

- number of TCP connections
  - Web1.1: **2**
  - BitTorrent: **~100** observed active
    - varies widely depending on
      - no. of torrents per user
      - maturity of swarm
      - config'd parameters

details suppressed:

- utilisation never 100%
  - but near enough during peak period
- on DSL, upstream constrains most p2p apps
  - other access (fixed & wireless) more symmetric

typical p2p file-sharing apps

- 105-114 active TCP connections altogether

1 of 3 torrents shown
- ~45 TCPs per torrent
- but ~40/torrent active

environment
Azureus BitTorrent app
ADSL+ 448kb upstream
OS: Windows XP Pro SP2

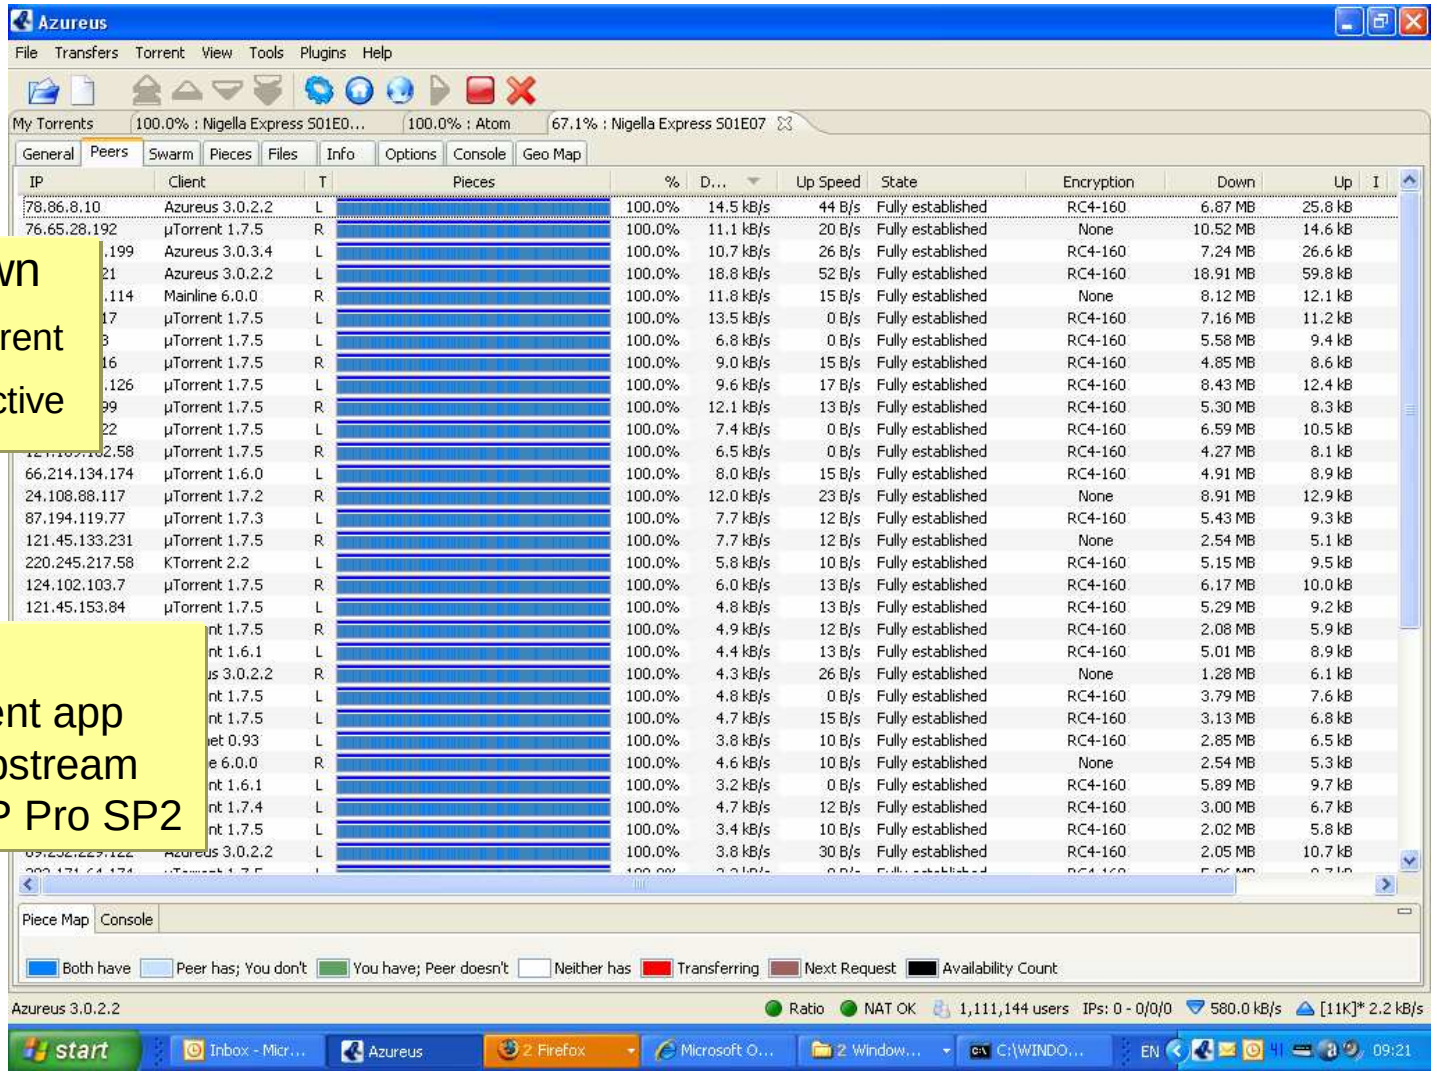

most users hardly benefit from bottleneck upgrade

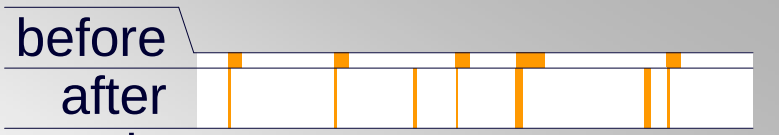

data limited flows
want rate more than volume

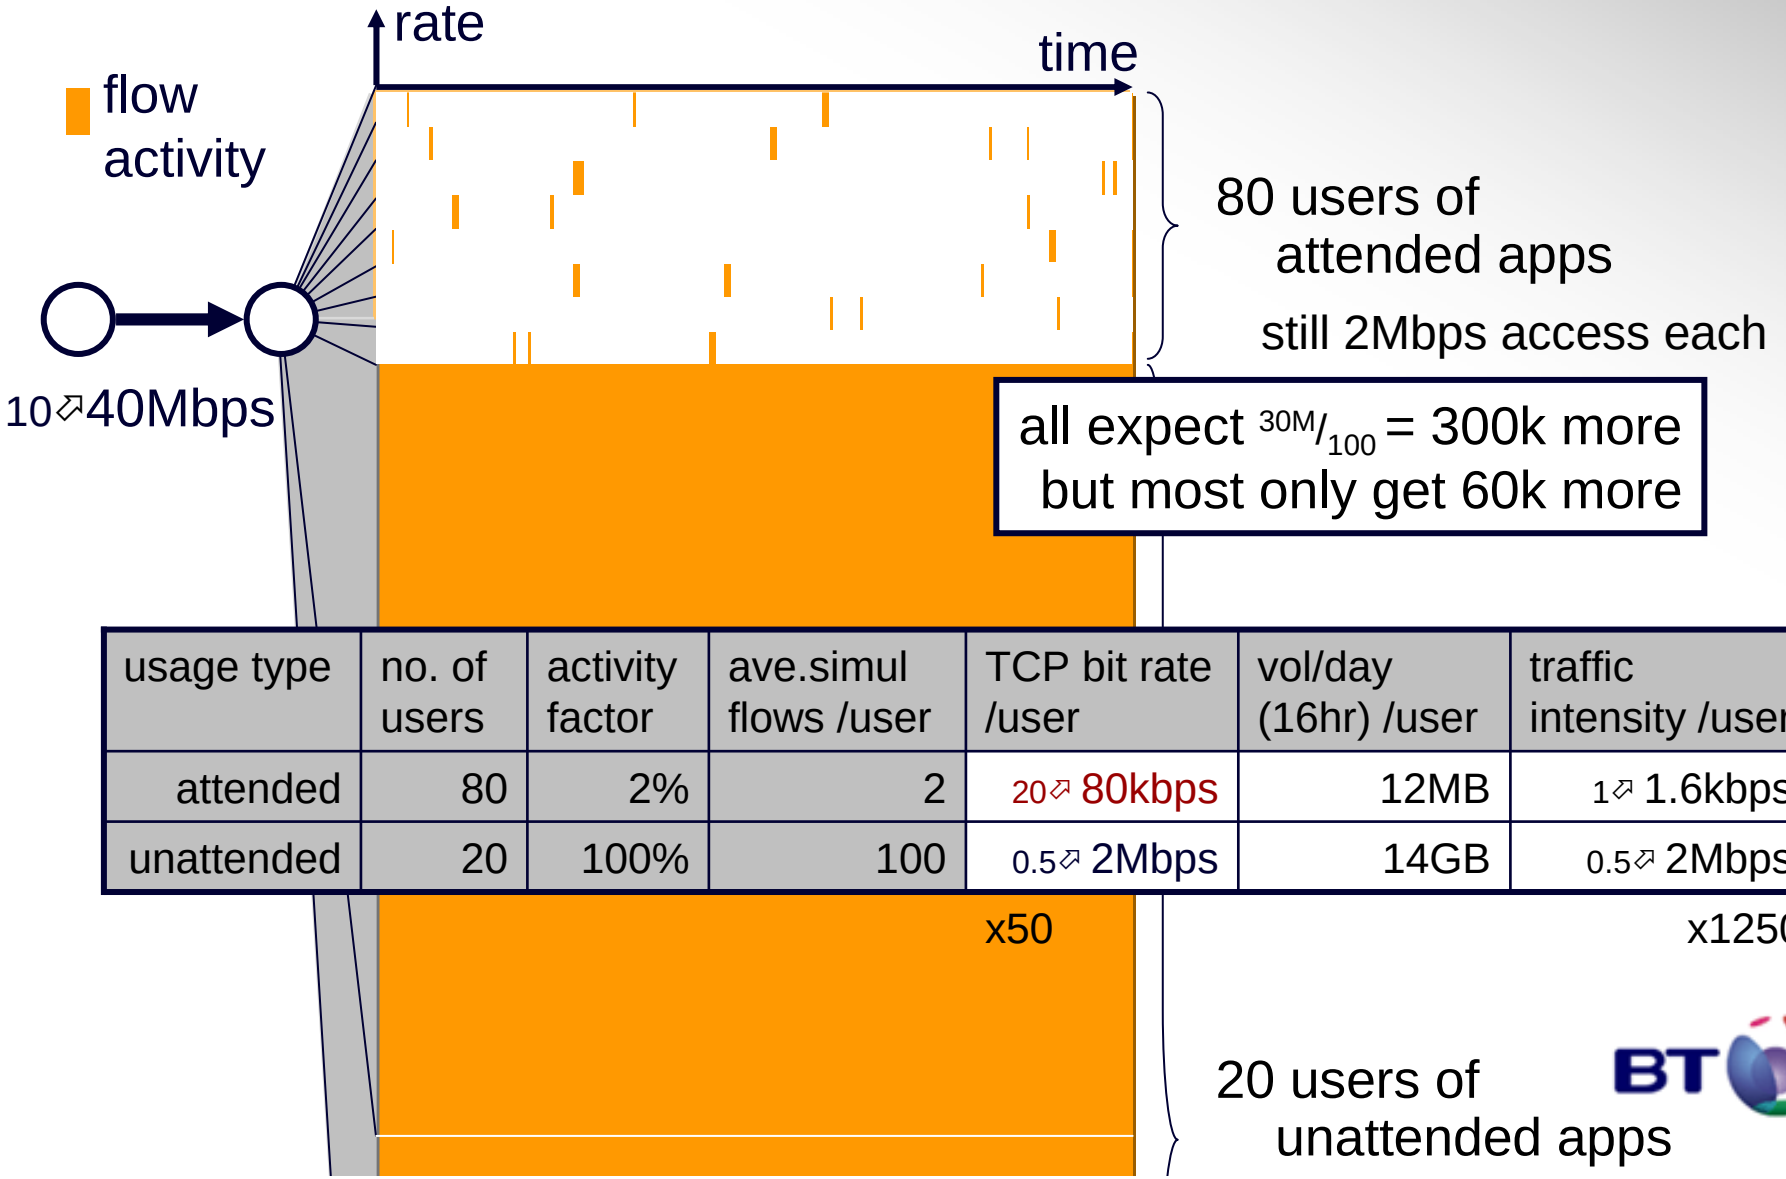

usage type	no. of users	activity factor	ave.simul flows /user	TCP bit rate /user	vol/day (16hr) /user	traffic intensity /user
attended	80	2%	2	20⇗ 80kbps	12MB	1⇗ 1.6kbps
unattended	20	100%	100	0.5⇗ 2Mbps	14GB	0.5⇗ 2Mbps

so what?

- fairness can’t be such a problem, the Internet works
  - we all have enough most of the time, even if A has more than B
  - Internet technical community likes to think this is due to its protocols
  - next few slides cast doubt on this complacency

concrete consequence of unfairness #1

higher investment risk

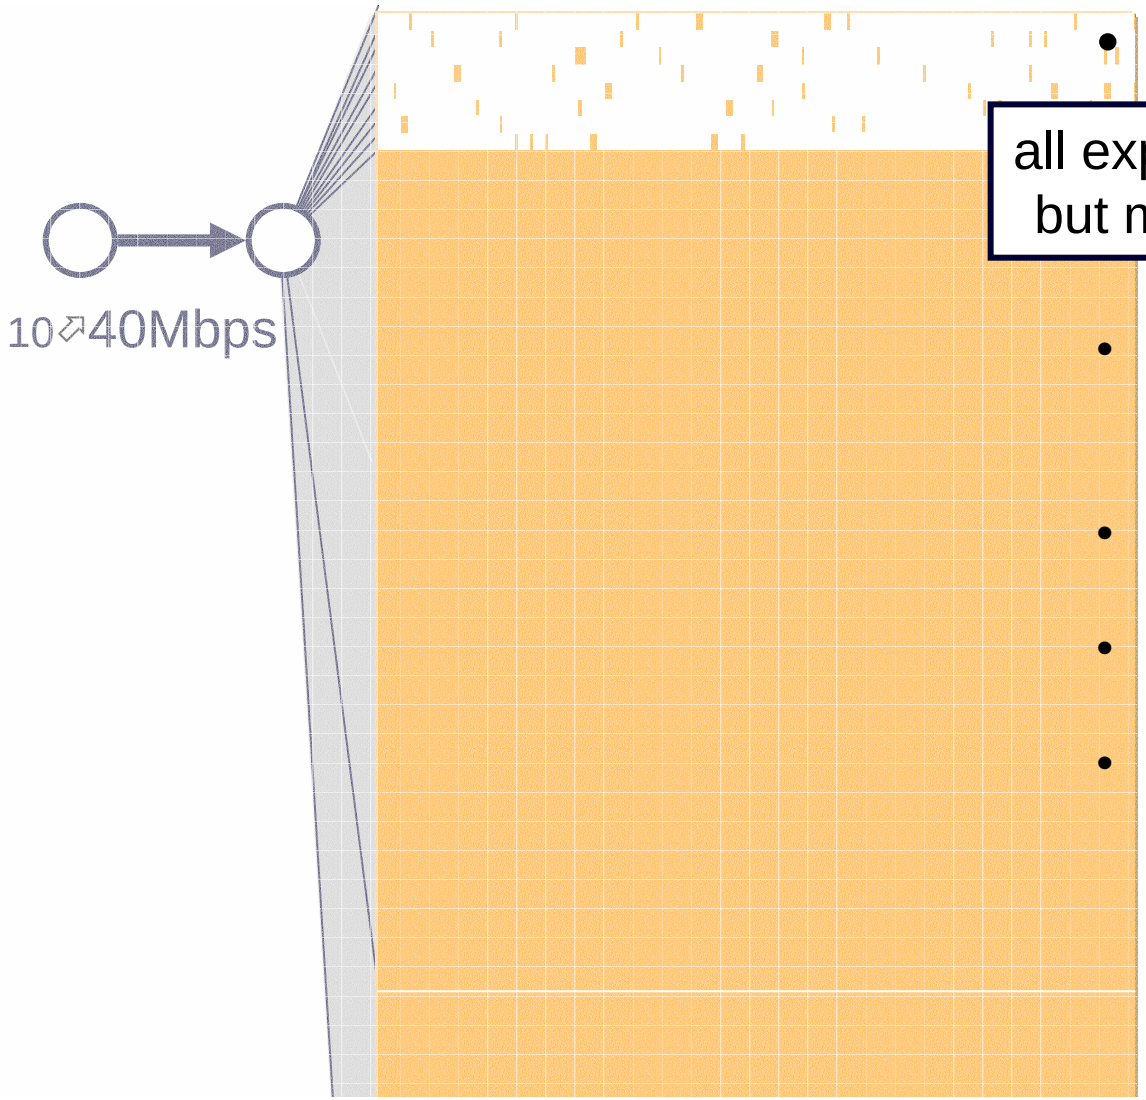

- recall

all expect $^{30M}/_{100}$ = 300k more
but most only get 60k more

- but ISP needs everyone to pay for 300k more
- if most users unhappy with ISP A's upgrade
- they will drift to ISP B who doesn't invest
- competitive ISPs will stop investing...

...but we still see enough investment

- main reasons
  - subsidies (e.g. Far East)
    - light users get ‘enough’ if more investment than they pay for
  - weak competition (e.g. US)
    - operators still investing because customers will cover the costs
  - throttling heavy users at peak times (e.g. Europe)
    - overriding TCP’s rate allocation

concrete consequence of unfairness #2

trend towards bulk enforcement

- as access rates increase
  - attended apps leave access unused more of the time
  - anyone might as well fill the rest of their own access capacity
- operator choices:
  a) either continue to provision sufficiently excessive shared capacity
  b) or enforce tiered volume limits

---

see CFP white paper “Broadband Incentives”

so the Internet way was wrong

and the operators were right?

- no, both were part right, part wrong

	the Internet way (TCP)	operators (& users)
degree of freedom	'flow rate equality'	'volume accounting'
multiple flows	✗	✓
activity factor	✗	✓
application control	✓	✗
congestion variation*	✓	✗

*another story

- both sides are failing to understand the strengths of the other

concrete consequence of unfairness #3

networks making choices for users

- characterisation as two user communities over-simplistic
  - heavy *users* mix heavy and light *usage*
- two enforcement choices

  a) bulk: network throttles all a heavy user's traffic indiscriminately
  - encourages the user to self-throttle least valued traffic
  - but many users have neither the software nor the expertise

  b) selective: network *infers* what the user would do
  - using deep packet inspection (DPI) and/or addresses to identify apps
- even if DPI intentions honourable
  - confusable with attempts to discriminate against certain apps
  - user's priorities are task-specific, not app-specific
  - customers understandably get upset when ISP guesses wrongly

BT

there are better solutions than fighting

think on this

- are these marketing spin for the same thing?
  a) slowing down heavy users
  b) allowing light users to go faster

- light usage can go much faster
  without appreciably affecting completion times of heavy usage

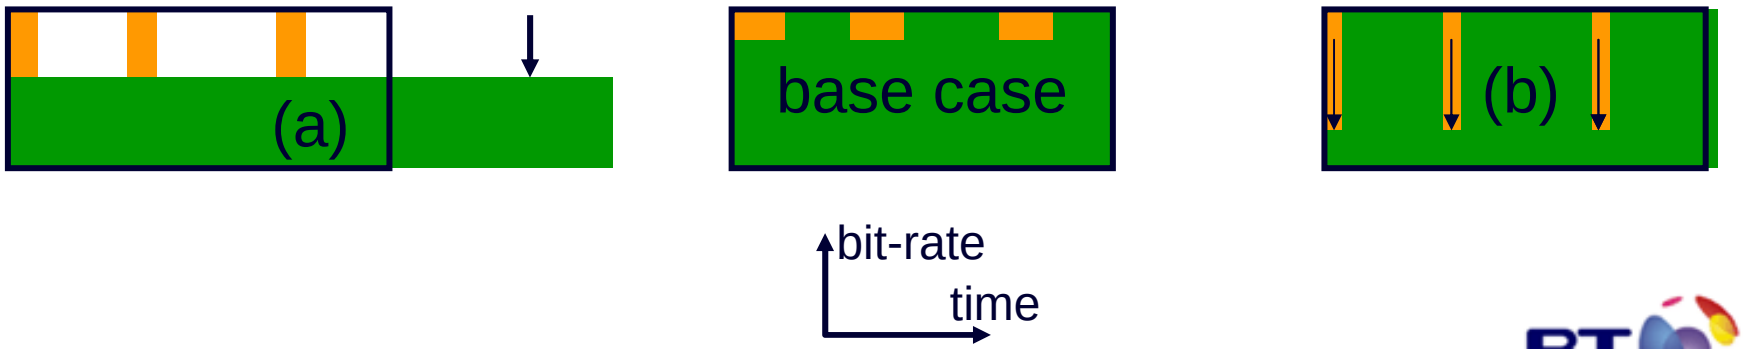

BT's two solutions (each yet another story)

- tactical (operational architecture)
  - "long term fair queuing"
- strategic (future Internet arch)
  - bulk edge congestion policing using "re-feedback"
  - encourages evolution of weighted TCP

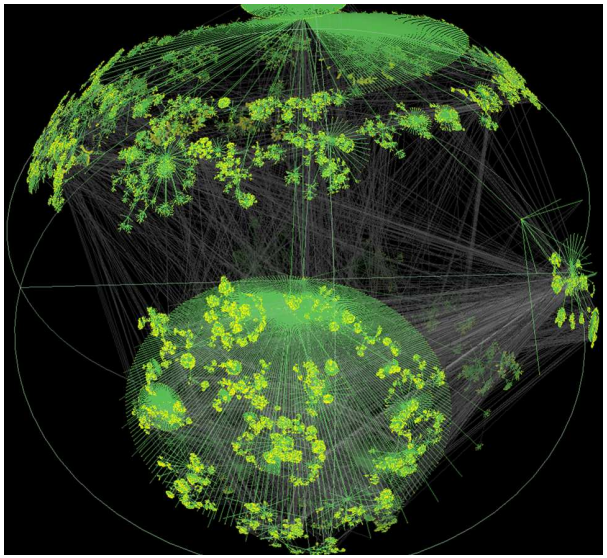

anyone will (still) be able
to use any link on the Internet
...without asking

whether NGN, cellular, ad hoc wireless,
public Internet, satellite, cable...

BT

Further reading

Problem Statement: We [the IETF] don't have to do fairness ourselves <www.cs.ucl.ac.uk/staff/B.Briscoe/projects/refb/#relax-fairness>

Q&A

freedom to limit the freedom of others