we don’t have to decide fairness ourselves

<draft-briscoe-tsvwg-relax-fairness-00.txt>

intent: build consensus then Informational

Bob Briscoe
Chief Researcher, BT
Toby Moncaster & Lou Burness
IETF-70 tsvarea Dec 2007

shifting IETF focus from fairness to accountability

	design-time	run-time
problem	IETF doesn't, can't and shouldn't decide fairness	
solution process	IETF's role: enable accountability for congestion	users, apps & operators can (optionally) make principled fairness choices
	IETF/IRTF can truly meet dynamic app req's and minimise congestion	
	best metric: congestion volume	

this talk primarily about the technical problem

- fairness is run-time, IETF is design-time

fair bottleneck bit-rate?

two incompatible partial worldviews

‘flow rate equality’	‘volume accounting’
per flow	per user
instantaneous	over time

- IETF aware that fairness should be per user
  - per flow is reasonable approx’n if users open similar no’s of flows

volume over peak period [B/hr]
instantaneous equivalent: traffic intensity [b/s]
= (bit rate when active [b/s]) x (activity factor [%])

base example

different activity factors

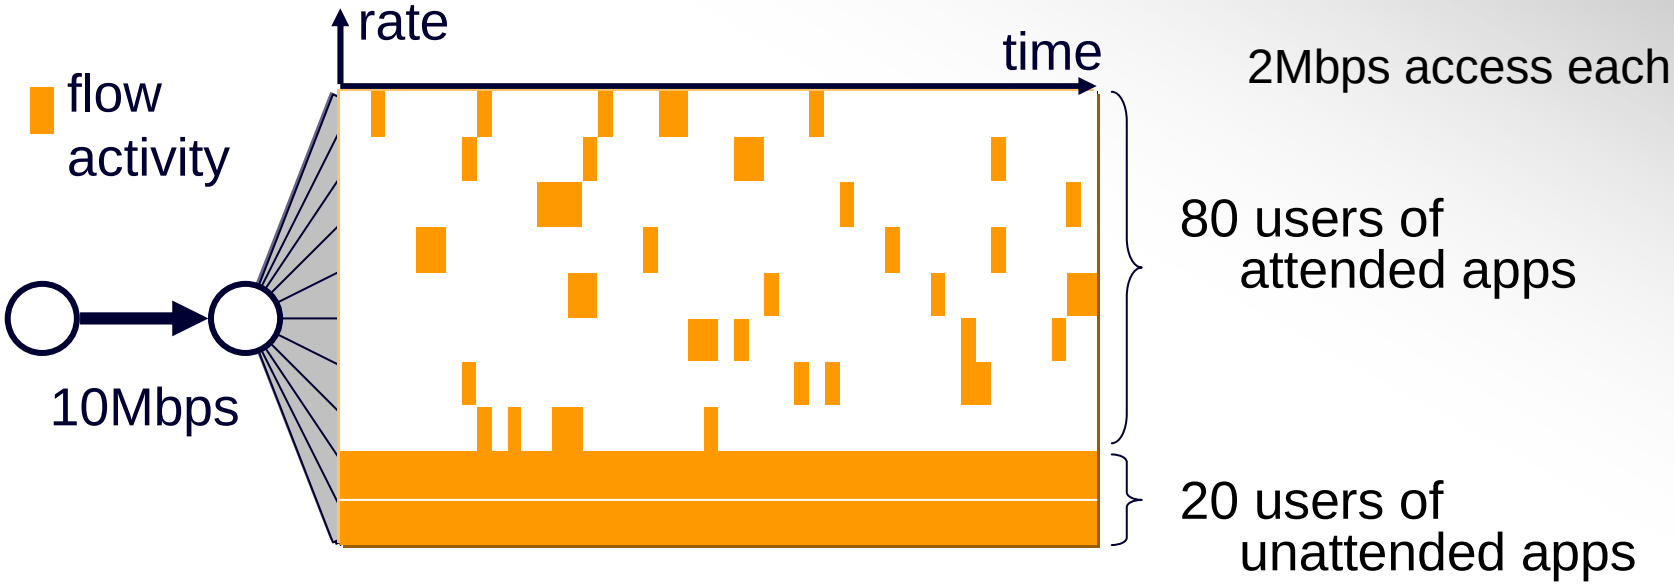

usage type	no. of users	activity factor	ave.simul flows /user	TCP bit rate /user	vol/day (16hr) /user	traffic intensity /user
attended	80	10%	=	357kbps	257MB	35.7kbps
unattended	20	100%	=	357kbps	2570MB	357kbps
				x1	x10	x10

realistic numbers?

there are elephants in the room

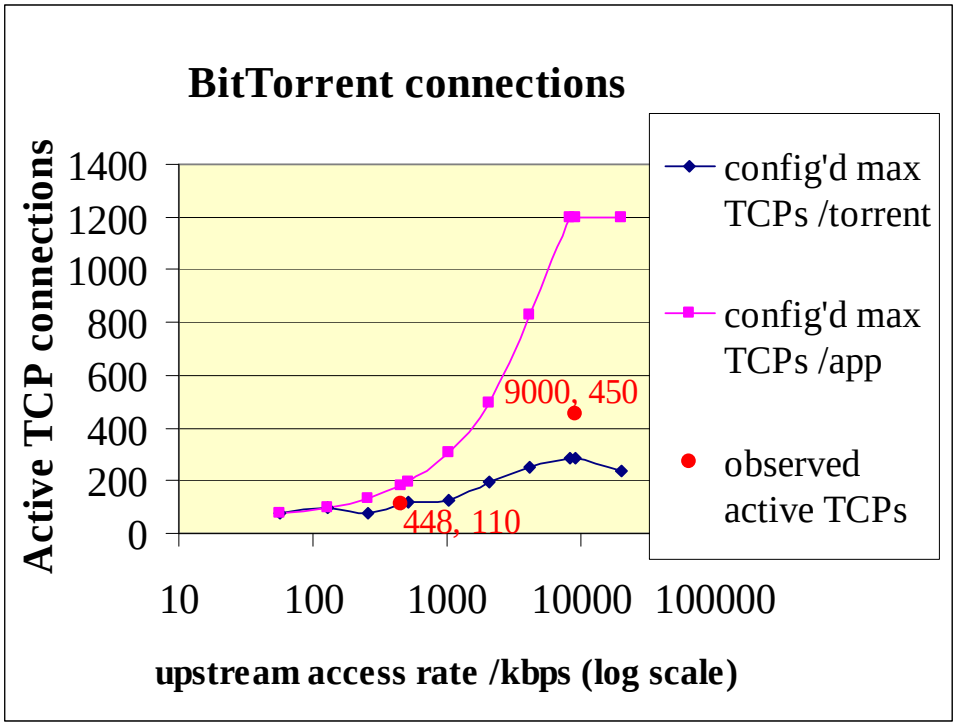

- number of TCP connections
  - Web1.1: **2**
  - BitTorrent: **~100**; see graph

details suppressed:

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

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

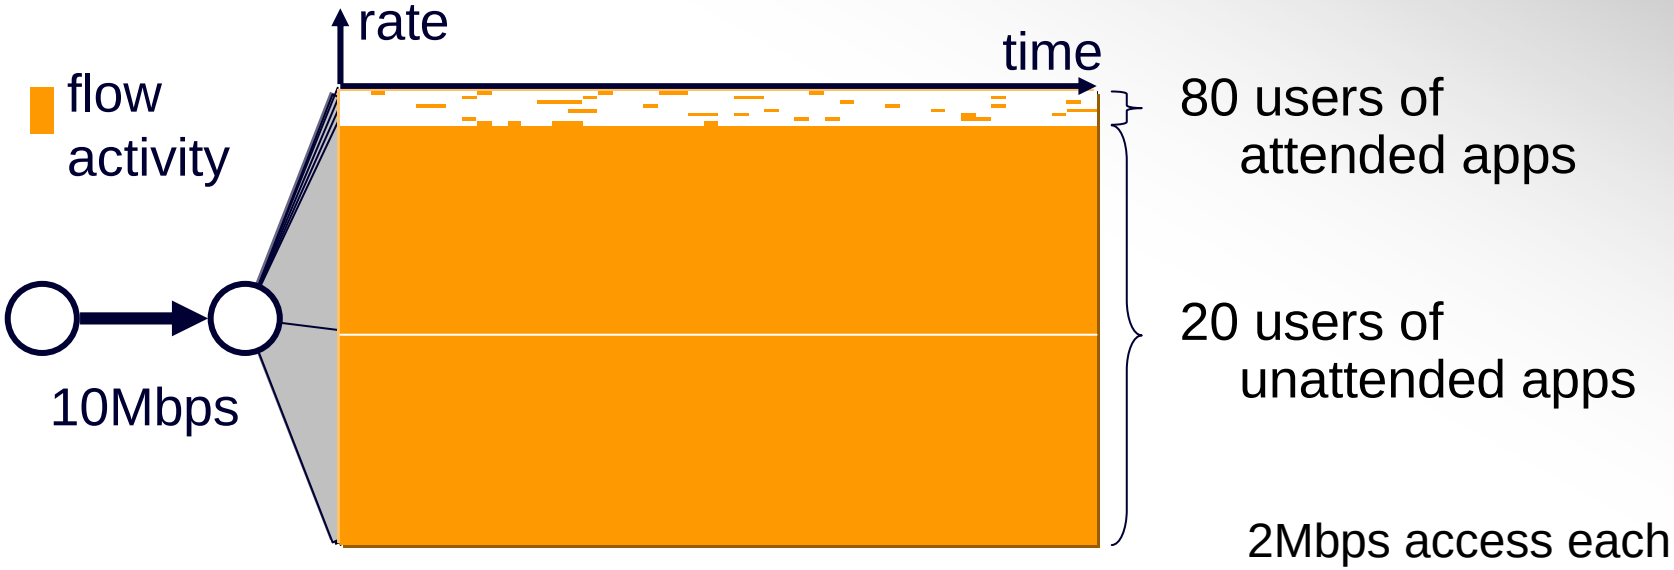

usage type	no. of users	activity factor	ave.simul flows /user	TCP bit rate /user	vol/day (16hr) /user	traffic intensity /user
attended	80	10%	2	10kbps	7.1MB	1kbps
unattended	20	100%	100	500kbps	3.6GB	500kbps
				x50	x500	x500

most users hardly benefit from bottleneck upgrade

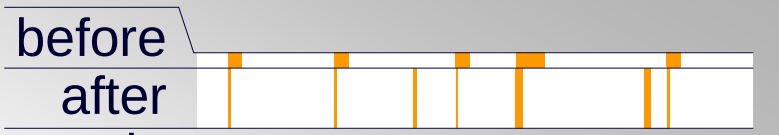

data limited flows
want rate more than volume

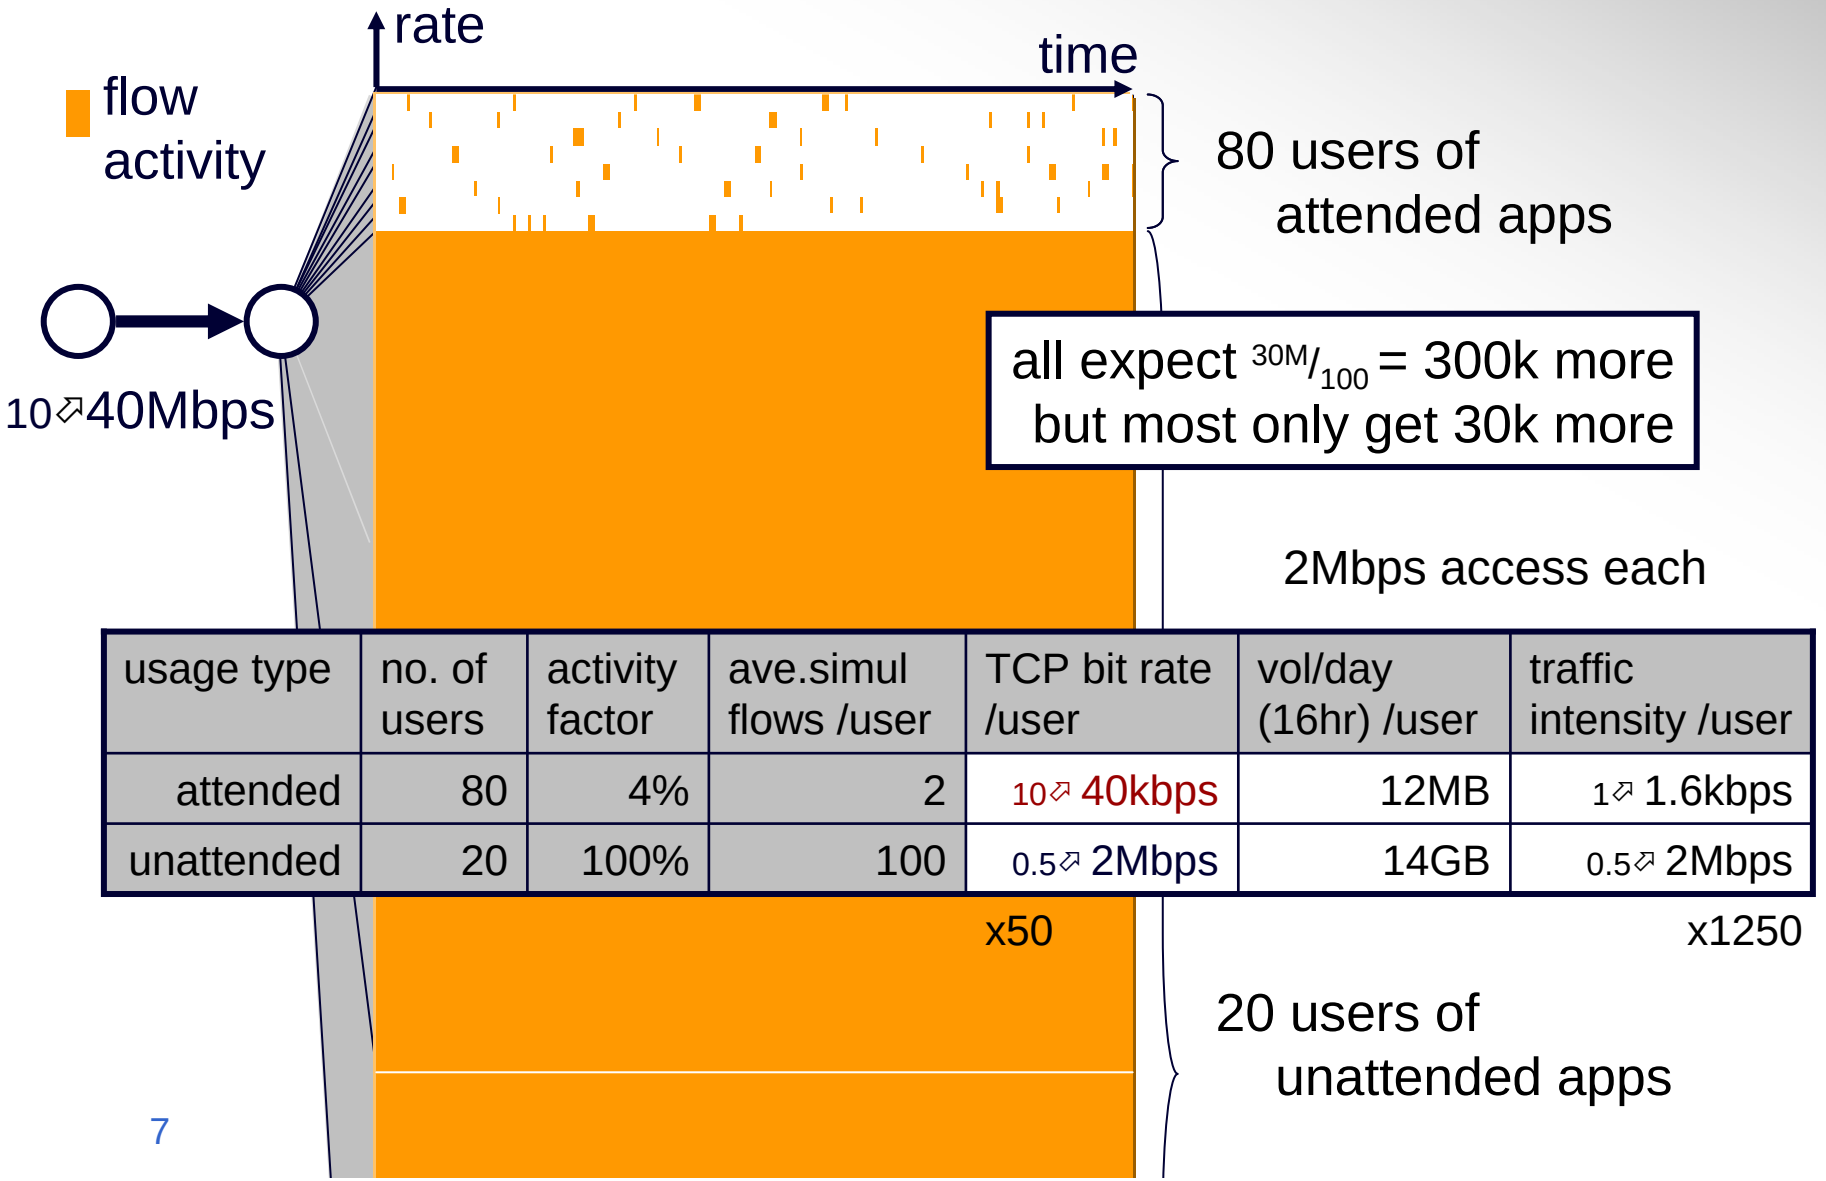

usage type	no. of users	activity factor	ave.simul flows /user	TCP bit rate /user	vol/day (16hr) /user	traffic intensity /user
attended	80	4%	2	10⇗ 40kbps	12MB	1⇗ 1.6kbps
unattended	20	100%	100	0.5⇗ 2Mbps	14GB	0.5⇗ 2Mbps

x50 x1250

20 users of unattended apps

volume accounting isn't the answer either

- fairer if heavy users get less bottleneck flow rate than light users
  - but heavy & light only defined by volume during 'the peak period'
  - effectively treats congestion very vaguely as
    - 0 everywhere off-peak
    - 1 everywhere on-peak
  - blind to whether the same volume causes extreme congestion or none

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

- message so far: 2 worldviews both claim same goal (fairness)
  - each strong over part of the problem space
  - but incompatible: one wants equal, the other wants unequal flow rates

enforcement of either goal is a separate issue (see later)

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
  - we like to think this is due to IETF protocols
  - next few slides cast doubt on this complacency

concrete consequence of unfairness #1

higher investment risk

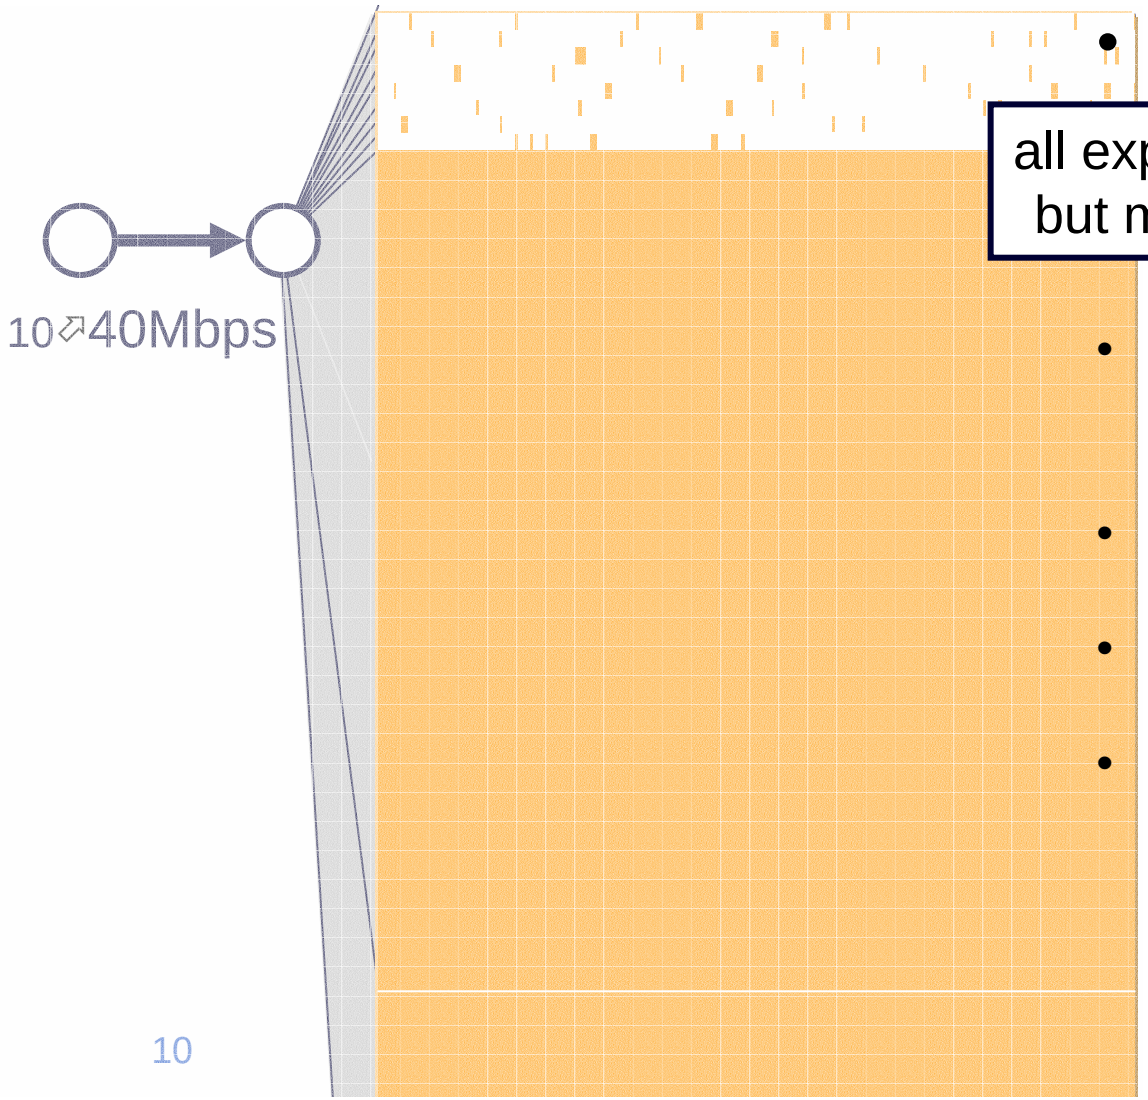

- recall

all expect $^{30M}/_{100}$ = 300k more
but most only get 30k 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 introduce tiered volume limits etc
- IETF needs to recognise & address the implications
  - bulk policing prevalent in best efforts architecture (cf. Diffserv)
  - e.g. should we distinguish a policer drop from a congestion drop?

concrete consequence of unfairness #3

networks making choices for users

- networks hit a problem once they start throttling
  - they could throttle 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
- many networks *infer* what the user would do
  - using deep packet inspection (DPI) to identify apps
- even if 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
- IETF needs to recognise & address the underlying need here
  - feature creep into network slows innovation (e2e principle)
  - better ways to fit traffic within limits (e.g. user/app-controlled endpoint s/w)

the problem

- IETF doesn't really decide fairness
  - whatever protocols *designed* to do, they are being *used* unfairly
- IETF can't really decide fairness
  - design-time body can't control run-time degrees of freedom
- IETF shouldn't decide fairness
  - shouldn't prejudge fair-use policy agreed between user & ISP
    - whether TCP, max-min, proportional or cost fairness

what does the IETF need to do?

- average rates – a run-time issue
  - introduce congestion accountability framework*****
  - give principled effective fairness control to users, apps & operators
  - offer an evolvable alternative to current kludges (DPI)
  - coexist with null enforcement
- transport dynamics – the design-time issue
  - IETF/IRTF protocols can truly satisfy dynamic application requirements while minimising congestion
  - rather than not really meeting app reqs, by being over-constrained

---

* TBA (Lou Burness +)

working towards BoF, not just about fairness, but also congestion collapse & DDoS

re-ECN / re-feedback one proposed solution

relaxing our transport design constraints

- currently we are trying to satisfy demanding app reqs
  - constrained by staying not 'much' more demanding than TCP
  - resulting protocols are 'over-constrained' and not app-developer's first choice
- once the big *average* rate fairness trade-offs move to run-time
- IETF/IRTF can judge which proposed transports better trade-off:
  - achieving the task effectively and
  - minimising unnecessary congestion to others during *dynamics*

- focus on the demanding *dynamics* questions:
  - when is a fast start fast enough? or too fast?
    [Limited slow start, etc]
  - how quickly should hi-speed transports allow in new flows?
    [HighSpeed TCP, FAST, etc]
  - how smooth can a transport be before it's effectively unresponsive?
    [TFRC, proprietary media players, etc]
  - what's the minimum congestion response of an aggregate?
    [PWE3, CAPWAP]

proposed core of solution

congestion harm metric

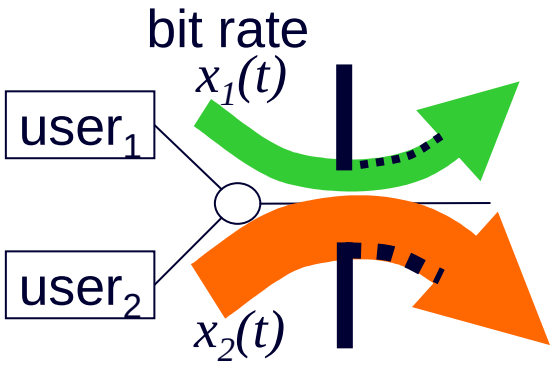

loss (marking) fraction $p(t)$

- partial insight from volume accounting
- but rather than only counting bytes during peak
  - count bit rate *weighted* by congestion, over time
  - result is easy to measure per flow or per user
    - volume of bytes discarded (or ECN marked)
- termed **congestion volume**

- a precise instantaneous measure of harm, counted over time
  - a measure for fairness over any timescale
  - and a precise measure of harm during dynamics

summary

shift IETF focus from fairness to accountability

	design-time	run-time
problem	IETF doesn't, can't and shouldn't decide fairness	users, apps & operators actually control fairness
	IETF's role: enable accountability for congestion	users, apps & operators can (optionally) make principled fairness choices
solution process	IETF/IRTF can truly meet dynamic app req's and minimise congestion	IETF protocols become first choice for demanding apps ☺
	best metric: congestion volume	

- problems will only get worse – driven by access rate increases

we don’t have to decide fairness ourselves

<draft-briscoe-tsvwg-relax-fairness-00.txt>

Q&A

context

3. a protocol solution: re-ECN <draft-briscoe-tsvwg-re-ecn-04.txt>
   - on hold while build consensus on the problem & requirements
   - other solutions welcome
0. dismantling flow rate fairness <draft-briscoe-tsvarea-fairness-02.pdf>
   - too polemical for IETF consensus
   - let this draft die – archived on my Web site and ACM CCR paper
1. the problem <draft-briscoe-tsvwg-relax-fairness-00.txt>
   - IETF doesn't decide fairness – this talk
2. solution requirements <draft-burness-tsvwg-...>
   - TBA

not pushing technical solution(s) at steps 1 & 2

- aimed more towards a 'congestion accountability' BoF

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

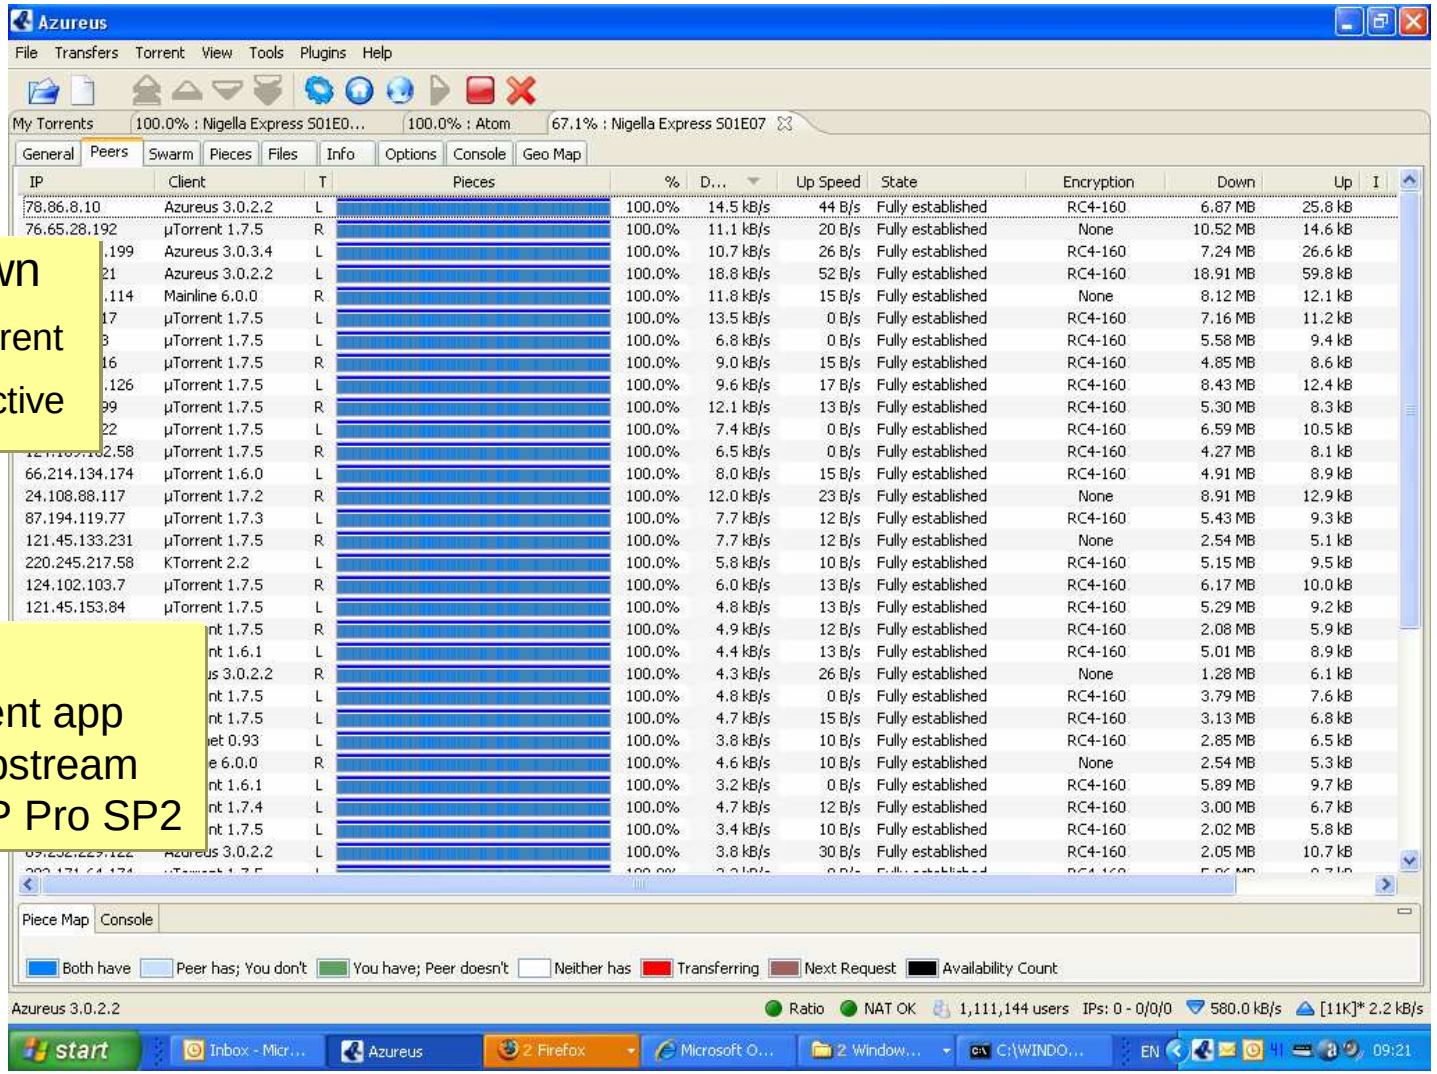

access growth just gets filled

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

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

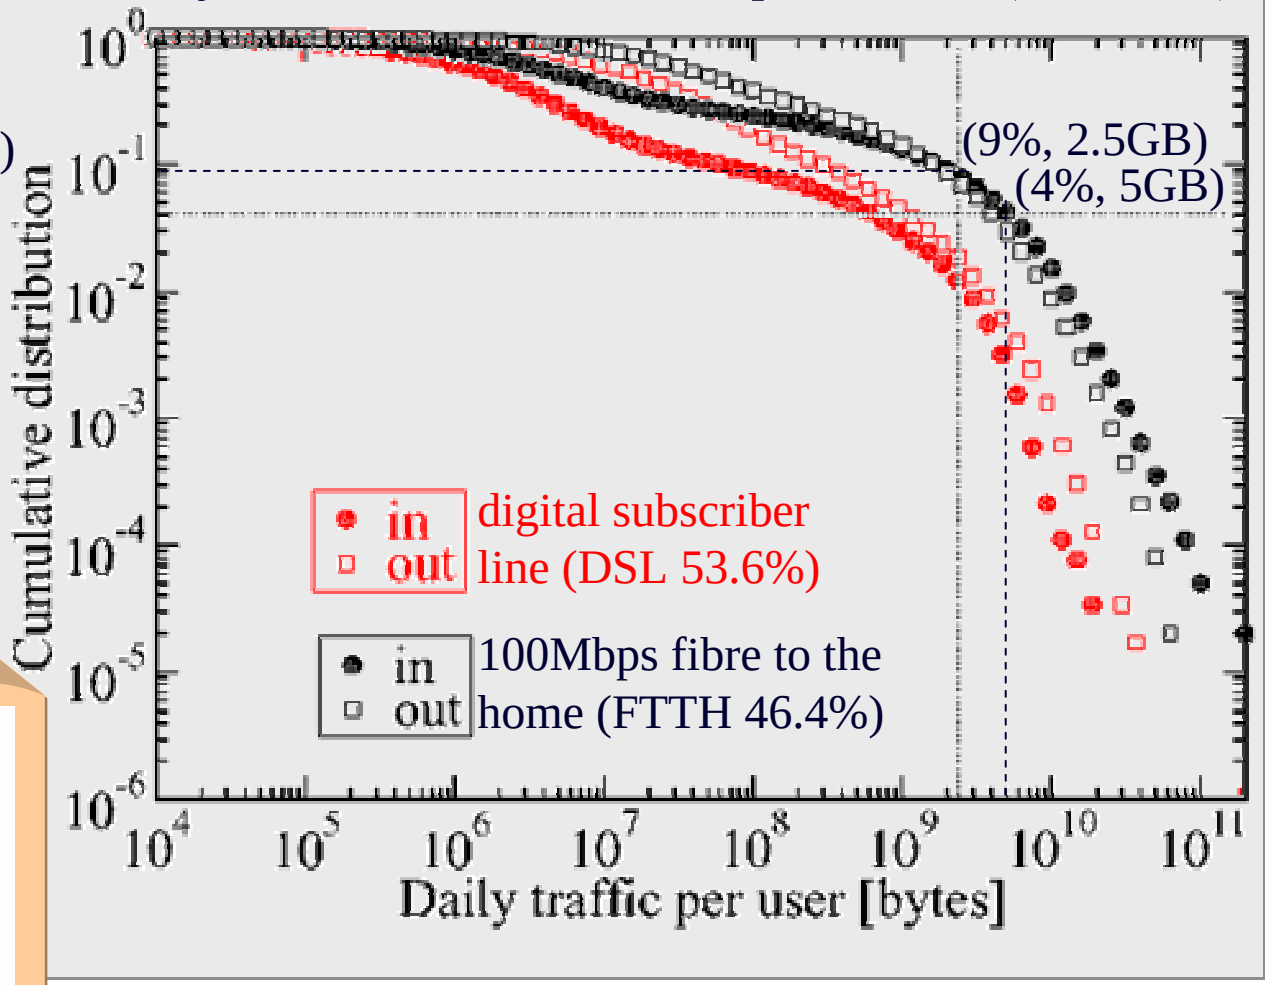

Changing technology shares of Japanese access market

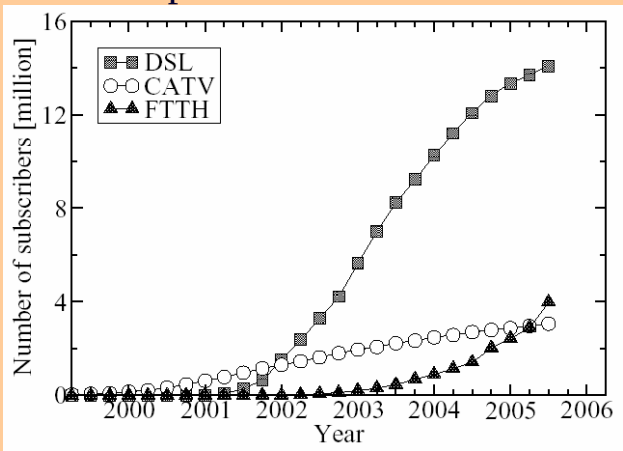

Courtesy of Kenjiro Cho et al
The Impact and Implications of the Growth in Residential User-to-User Traffic, SIGCOMM'06

concrete consequence of unfairness #4
starvation during anomalies & emergencies

- fairness concerns become acute during stress
  - more traffic or less capacity than expected
- if fairness decided at run-time
  - common policy probably ‘you get what you paid for’
- concern: unsavoury for emergencies
  - all flows should make some progress, not just the rich
- agree with concern, but current approach not right
  - video downloads get 50x rate of emergency messages?*
- policy decisions for users, ISPs, regulators, not IETF
  - e.g. ISP might freeze paying to override congestion limits

* Henchung earthquake, 26 Dec ’06, see I-D

accountability metric

congestion volume

precisely measures instantaneous harm from flow rate dynamics rather than just average flow rate