flow rate fairness dismantling a religion

<draft-briscoe-tsvarea-fair-02.pdf>

Bob Briscoe

Chief Researcher, BT

IETF-69 tsvwg Jul 2007

where we stand

task:
how to share all the parts
of a huge, multi-provider
packet multiplexer
between competing processes

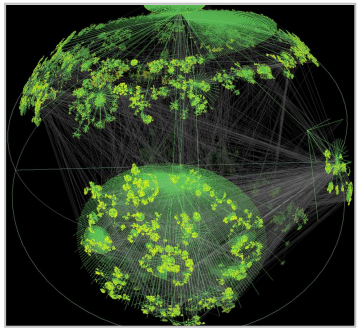

- equal flow rates at a bottleneck sufficed in the past
  - but not <u>just</u> a question of whether every flow makes progress
  - virtually any share ‘works’
- using equal rates between flows as the goal in the past
  - has caused apps to evolve that open more flows, for longer
  - >60-80% of traffic now from apps opening numerous, very long flows
- that’s cool
- but it actually just says
  - “If you take more, you get more”
  - it shows our protocols don’t affect fairness at all
  - because they missed the bigger picture
- we need to admit this

what should the IETF do?

- we've got nothing to stop much more selfish apps evolving
  - DDoS already with us, p2p interactive video is growing
- today fairness enforcement all outside IETF
  - kludged, complex or freezes-in today's apps
  - deep pkt inspection, bottleneck policers,
    volume caps, volume pricing, emailed warnings
- IETF goal #1?
  - simplest possible effective fairness enforcement, but embrace diversity
    - cellular, NGN, ad hoc wifi, campus, corporate, public
  - to replace current kludges with evolvable alternatives
  - and protect against possible future fairness problems
  - and coexist with null enforcement

don’t mistake ‘add’ for ‘replace’

- “equal flow rates are ~~fair~~ useful” will be part of the future
  - not as a goal, just an allocation that ‘works’ and already exists

- IETF goal #2?
  - any ISP can choose not to deploy an enforcement mechanism
  - but its neighbours can choose to make it accountable for the effect on others

who should judge fairness?

- the existing way (equal flow rates)
  - the IETF broadly judges what's fair (but everyone can actually take more)
- new way (congestion volume metric)
  - IETF protocols enable local judgements of fairness
  - subsets can determine their own fairness policy (Universities, US, NATO)
  - globally, cost-fairness arbitrates between the subsets

- any fairness enforcement won't be as simple as none
  - we've suggested one pretty simple mechanism based on ECN (re-ECN)
    - yes, ECN is more complex than drop
  - simpler and more effective than current kludges
  - and application-neutral

updated 01⇨02 draft

- diffs and alt formats (courtesy of rfcdiff & xml2rfc tools) at: <http://www.cs.ucl.ac.uk/staff/B.Briscoe/pubs.html#rateFairDis>
- comments from presenting at IETF-68 tsvwg
  lots of (on & off list) email
- main changes from previous draft-01, clarifications
  - applicability within other Diffserv classes than BE
    - and for congestion of other lower layer resources (radio, battery etc)
  - we DO NOT recommend or require user congestion pricing (that's what we've solved)
  - we DO NOT recommend or require per flow policing (redundant with per-user policing)
  - cost fairness **<≠>** re-ECN
  - why congestion volume is so important (considerably clarified)
- this draft will now die (archived at above URL and ACM CCR paper == -00)
  - any parts of the text of this I-D are available for copy & paste to other I-Ds
  - Lou Burness volunteered to edit (+co-authors from list) a forward looking informational I-D

fairness metric
congestion volume

- flow rate has to be averaged and can't be integrated over time
- congestion volume is instantaneous and integrates over time

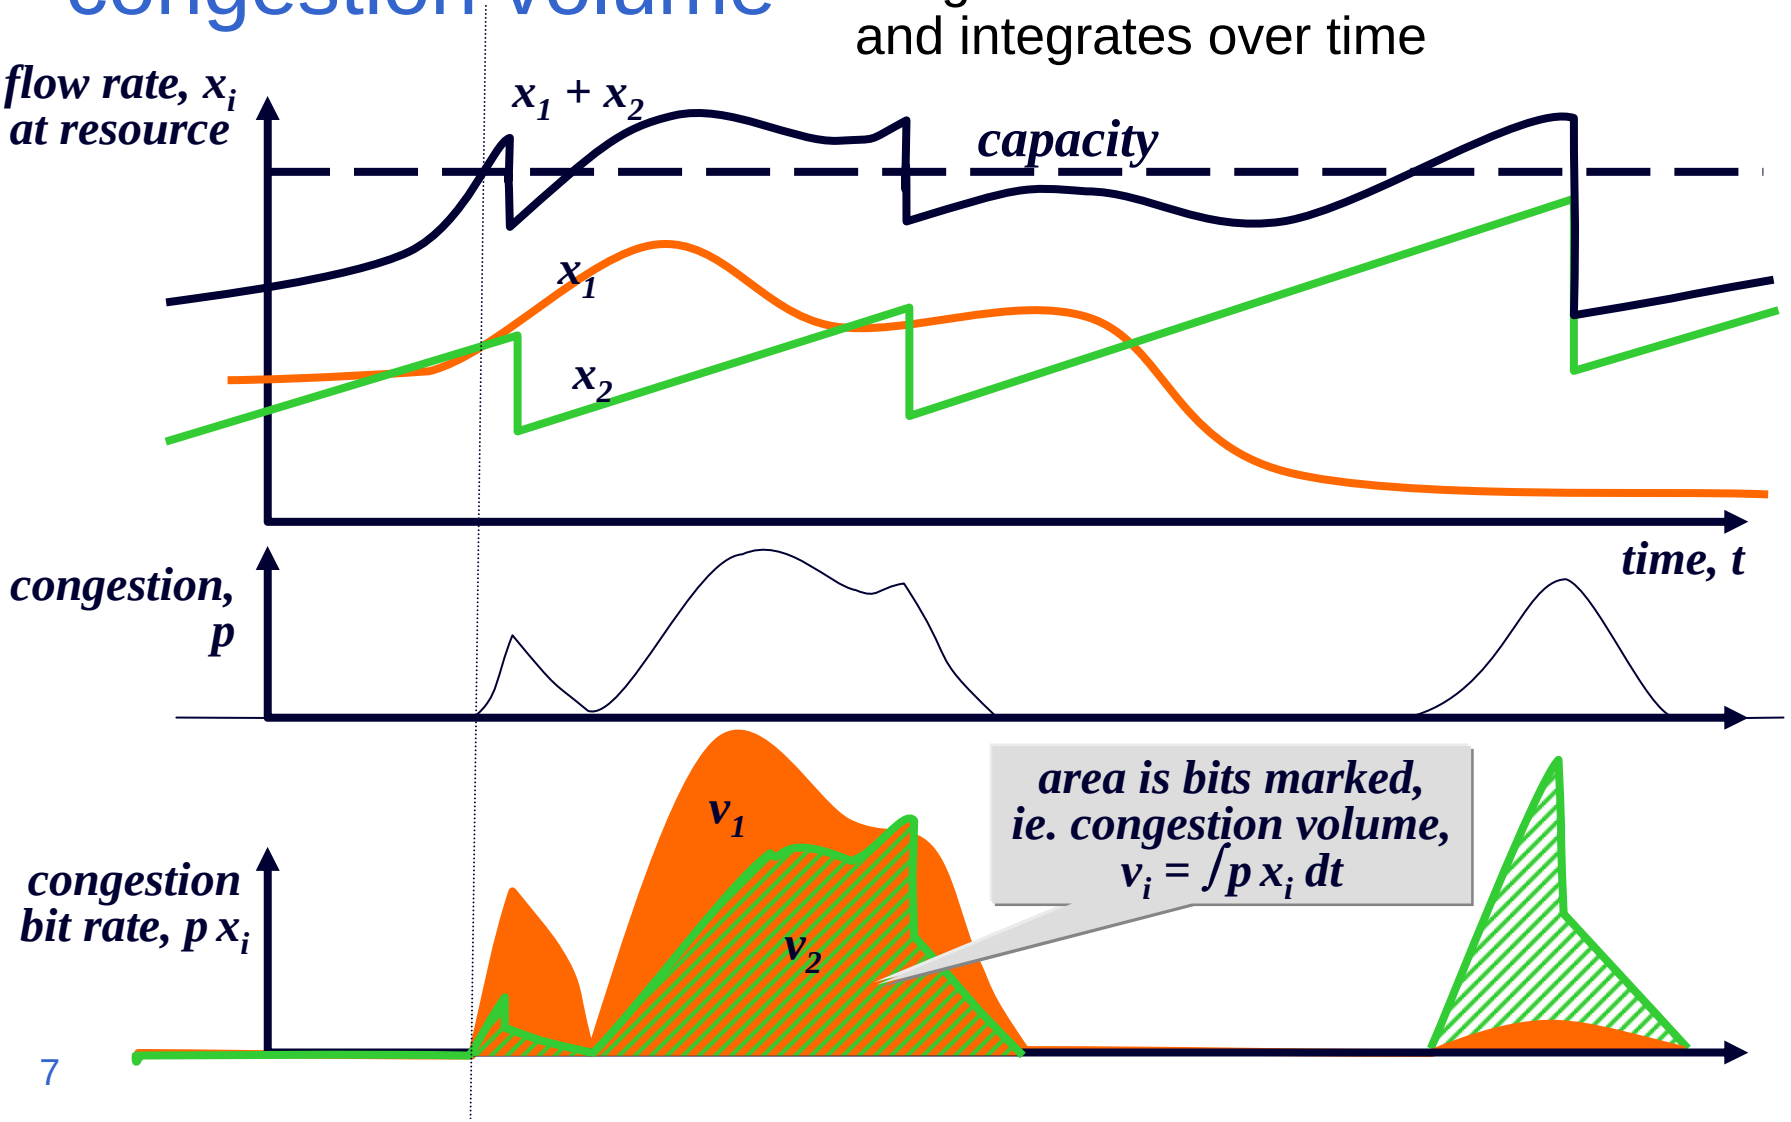

flow rate fairness
dismantling a religion
<draft-briscoe-tsvarea-fair-02.pdf>

Q&A

Bar BoF “re-ECN next steps”

Wed 25 July 1300-1500, Red Lacquar,
Palmer Ho Hilton, Chicago

background papers on re-ECN:

<http://www.cs.ucl.ac.uk/staff/B.Briscoe/projects/refb/>
including particularly
<draft-briscoe-tsvwg-re-ecn-tcp-04.txt>

calibrating ‘cost to other users’

- a monetary value can be put on ‘what you unsuccessfully tried to get’
  - the marginal cost of upgrading network equipment
    - so it wouldn’t have marked the volume it did
    - so your behaviour wouldn’t have affected others
- competitive market matches...
  - the cost of congestion volume
  - with the cost of alleviating it

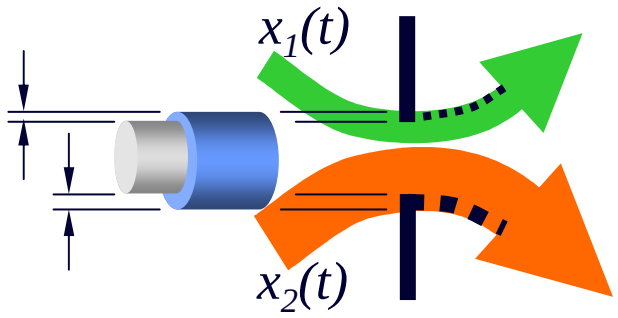

note: diagram is conceptual

congestion volume would be accumulated over time

capital cost of equipment would be depreciated over time

- congestion volume is not an extra cost
  - part of the flat charge we already pay
  - but we can’t measure who to blame for what
  - if we could, we *might* see pricing like this...

access link	congestion volume allow’ce	charge
100Mbps	50MB/month	€15/month
100Mbps	100MB/month	€20/month

- NOTE WELL
  - IETF provides the metric
  - industry does the business models