flow rate fairness dismantling a religion

<draft-briscoe-tsvarea-fair-00.pdf>

Bob Briscoe
Chief Researcher, BT Group
PFLDnet Feb 2007

today's shares are just the result of a brawl

- flow rate fairness is not even wrong
  - it doesn't even answer the right questions
  - it doesn't allocate the right thing
  - it doesn't allocate between the right entities
- how do you answer these questions?
  1) how many flows is it fair for an app to create?
  2) how fast should a brief flow go compared to a longer lasting one?

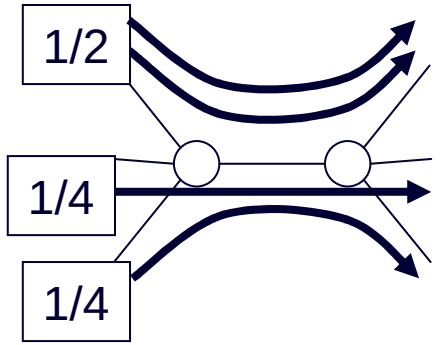

fair allocation... of what? among what?
of ‘cost’ among bits

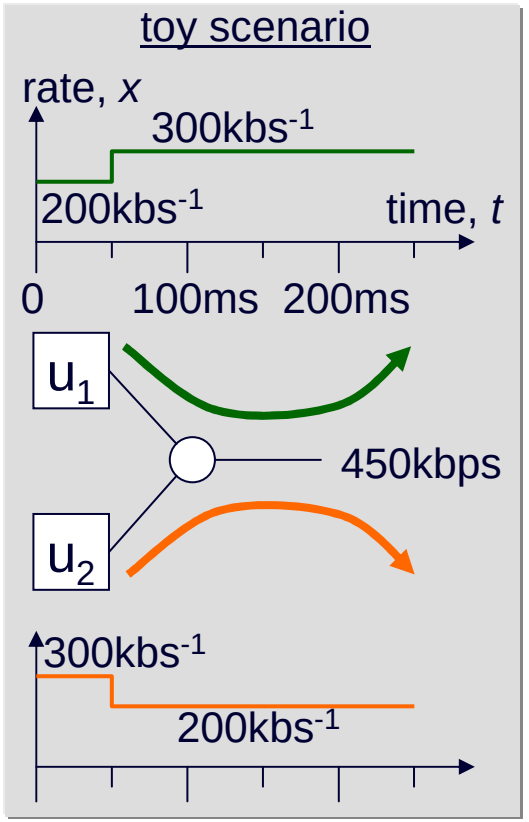

- cost of one user’s behaviour on other users
  - congestion volume = instantaneous congestion...
    - ...shared proportionately over each user’s bit rate
    - ...over time
- instantaneous congestion

  p = 10%
- congestion volume, $v = x(t).\Delta t.p(t)$

  v_1 = 200kbs^{-1} x 50ms x 10% + 300kbs^{-1} x 200ms x 10%

  = 1kb + 6kb = 7kb

  v_2 = 300kbs^{-1} x 50ms x 10% + 200kbs^{-1} x 200ms x 10%

  = 1.5kb + 4kb = 5.5kb
- as $\Delta t \rightarrow \delta t$, integrates easily & correctly over time and over flows
  - ≡ volume of data each user sent that was dropped (if loss-based)
  - ≡ volume of data each user sent that was congestion marked (if ECN-enabled)

toy scenario for illustration only; strictly...
- a super-linear marking algorithms to determine p is preferable for control stability
- the scenario assumes we’re starting with full buffers

enforcement of fairness

- if it's easy to 'cheat', it's hardly a useful fairness mechanism
  - whether intentionally or by innocent experimentation
- if every flow gets equal rate
  - the more flows you split your flow into, the more capacity you get
  - fairness per source-destination pair is no better
    - Web/e-mail hosting under one IP addr
    - stepping stone routing (cf bitTorrent)

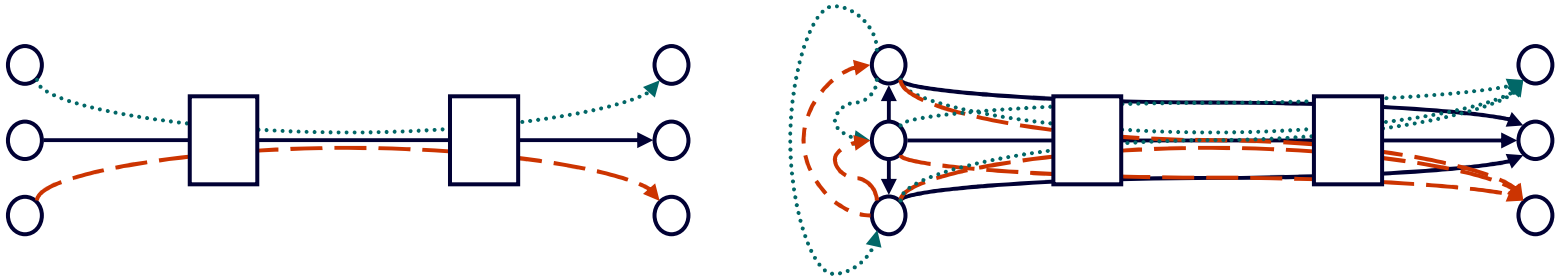

- by design, cost allocation among *bits* is immune to such cheating