



# Congestion Exposure

Bob Briscoe  
Chief Researcher, BT  
Jul 2010

This work is partly funded by Trilogy, a research project supported by the European Community  
[www.trilogy-project.org](http://www.trilogy-project.org)



## between a rock and a hard place

- proper transports can fill available capacity, but...
- what share should each get when they coincide?
- previous talks
  - economics says users would find answer themselves
  - if charged for their contribution to incipient congestion
- but unpredictability of congestion billing is unpopular
- consumers & businesses want flat fee
- network operators want engineered control
  - scary to depend on rational customers' price responses

## flat fee as if congestion charged

- we want apps to somehow behave as if the user is congestion charged, but without congestion charging
- need to allow network operators to set and enforce limits on each user's contribution to congestion
- “contribution to congestion” is congestion-volume
  - congestion-volume = volume x congestion (units of bytes)
  - congestion-rate = rate x congestion (units of bps)
  - e.g. 1Mbps flow x 0.1% congestion  
= 125 bytes congestion-volume in 1 second

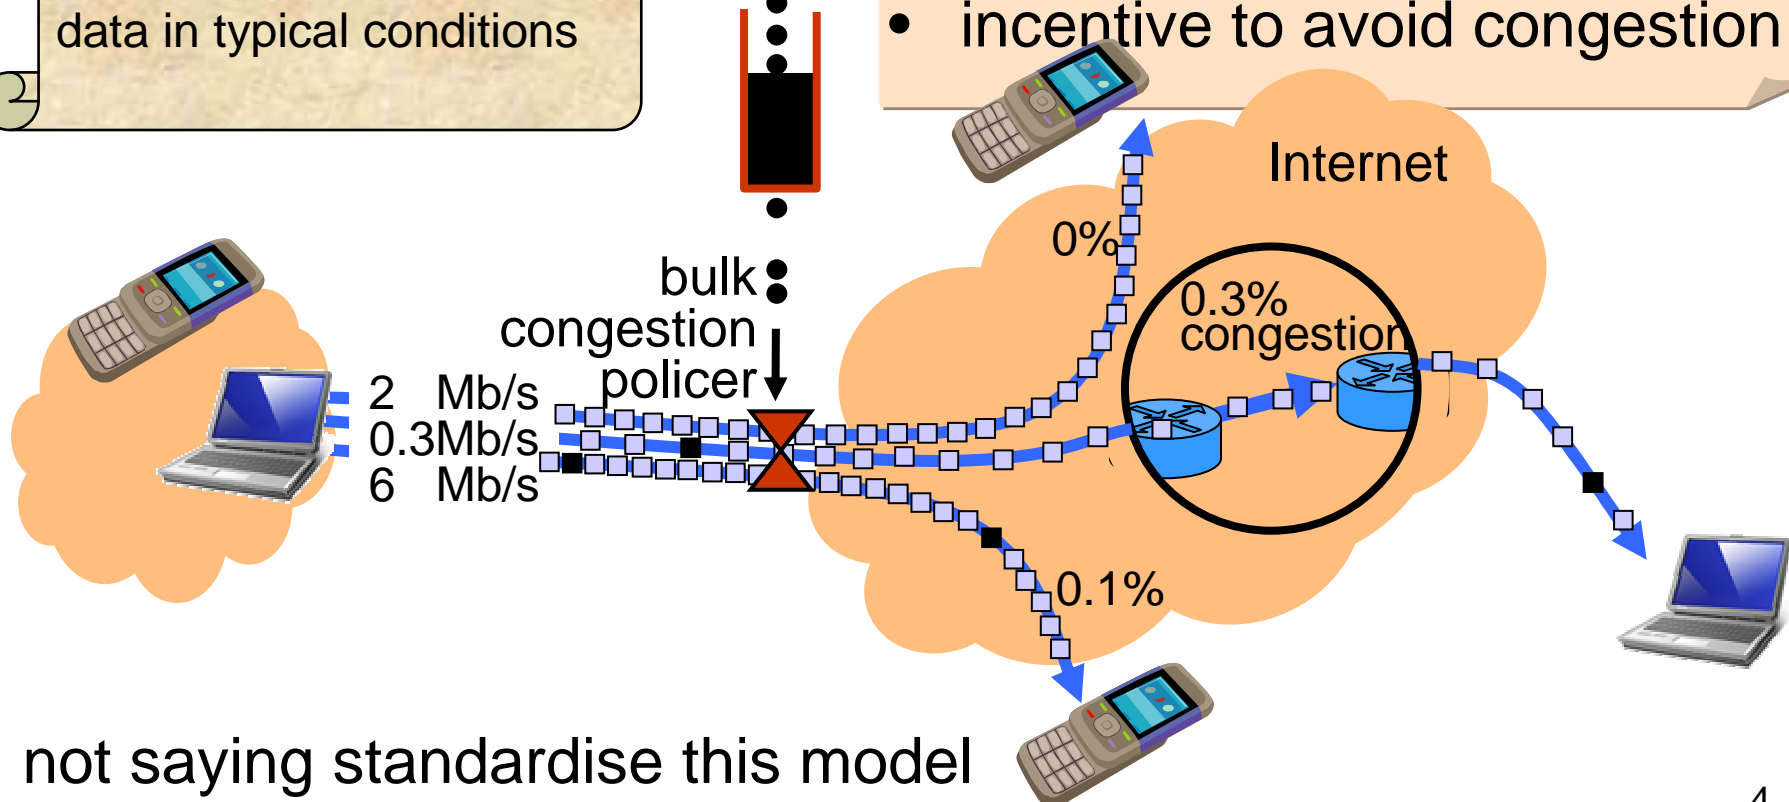
# example: flat fee congestion policing

## Acceptable Use Policy

'congestion-volume'  
allowance: 1GB/month

Allows ~70GB per day of  
data in typical conditions

- only throttles traffic when your contribution to any congestion in the Internet exceeds your allowance
- incentive to avoid congestion





not saying standardise this model  
example of what an operator should be able to do

# IETF task: Congestion Exposure (ConEx)



- but...
  - Internet architected for hosts to manage congestion
  - network can see utilisation, but not path congestion
- IETF task: provide feasible way for network operators to measure and control congestion-volume
  - needs to be as easy to measure as volume
  - and as transparent to verify and agree as volume
- Congestion Exposure (ConEx) working group
  - sender exposes expected congestion in IP header
  - IPv6 only initially and focus on partial deployment
- a generative technology: IETF merely defines the protocol
  - optional for networks and hosts
  - but networks can create incentives for sender to use it
    - and to be truthful
  - industry players and economics will drive how it is used

# what's wrong with TCP?

- surely TCP responds *as if* loss were a congestion charge?
- yes but... if you had to pay for congestion
  - you would weight each TCP very differently, not all the same



weighted TCP  
as if congestion  
charged



- problem:  
nothing to limit how much you use TCP
  - open more TCP sessions and you get more capacity
  - hand more data to TCP & it occupies capacity for longer
- anyway, using TCP is optional for an app


# what's wrong with current traffic controls?

- ISPs, enterprise, campus,... network operators
  - faced with competition, regulation, budget constraints
- currently some complement capacity investment with traffic controls
  - aiming to limit the most costly users
- economics says incremental cost of traffic = congestion
  - so don't traffic controls limit users contributing most congestion?
- Well, no... network cannot see congestion
- so networks limit what they can see...
  - instantaneous bit-rate, 95%ile, volume at peak time, p2p apps
  - piecemeal – when one doesn't work, try adding more...

outcome:

## an architectural soup of network controls

- traffic controls appear closer to ideal behaviour
- but with downsides
  - not user-controlled – they infer what the user wants
  - violate architectural coherence (e.g. DPI vs IPsec)
  - costly to manage complexity & unpredictable behaviour





## summary

- without Congestion Exposure, the Internet is far from working “as if there was congestion charging”
- no wonder the net neutrality debate is so confused
  - both host control & network control are severely lacking
- can't have flat fee as if congestion charging
- can't limit user's contribution to congestion
- network cannot see congestion
- fixing this is the Congestion Exposure (ConEx) goal

## more info...

- The whole story in 7 pages

- Bob Briscoe, "[Internet Fairer is Faster](http://bobbriscoe.net/projects/refb/#fairfastWP)", White Paper (Jun 2009)  
<<http://bobbriscoe.net/projects/refb/#fairfastWP>>

available from the re-feedback project page:

<<http://bobbriscoe.net/projects/refb/>>

<[bob.briscoe@bt.com](mailto:bob.briscoe@bt.com)>

- ConEx IETF working-group

<<http://datatracker.ietf.org/wg/conex/charter/>>

<[conex@ietf.org](mailto:conex@ietf.org)>

# Congestion Exposure


Q&A...



& spare slides...

## something like LEDBAT?

- surely LEDBAT behaves like this?



- but current traffic management discourages LEDBAT
  - LEDBAT still transfers high volumes, so is still targeted
  - LEDBAT used for applications like P2P, so is still targeted
  - LEDBAT is prevented from working by 'fair' queuing
- so LEDBAT focuses on the home gateway queue
  - hard to help other users when the ISP cannot tell :(