interconnect QoS settlements & impairments

Bob Briscoe
BT Group CTO

Q1. technical capabilities needed to support acceptable revenue models for providers?

what are providers trying to achieve?

- selling QoS = managing risk of congestion
  1. ranking demand so insufficient willingness-to-pay self-rejects
  2. and/or exploiting a monopoly position (perhaps only over a route)

1. push-back from congestion only requires congestion charging
   – peak-demand and volume charging are imperfect but pragmatic proxies
2. exploiting monopoly could require any sort of charging model
   – but must still push-back from congestion at some timescale

- a game is playing out, converging on near-perfect competition
  – play the game conceptually and deploy the end-game (congestion pricing)?
  – or play the game out in full? deploying/withdrawing many models on the way

Network cost economics (not market pricing) (perfect competition)

➔ infrastructure cost is sunk	➔ installation fee
➔ operational costs are usage independent	➔ monthly fee
➔ usage and congestion cost **operator** nothing	➔ 0
➔ congestion damages service to **user**	➔ congestion pricing
➔ congestion income pays for infrastructure upgrade	➔ installation fee ➔ 0

sender or receiver pays? recap

- two part tariff
  - sending domain pays $C = \eta X + \lambda Q$ to r'cving domain per accounting period
  - X is capacity @ price η
  - Q is QoS/usage-related (volume, peak demand, congestion) @ price λ
  - both prices relatively fixed
- usage related price $\lambda \geq 0$ (safe against 'denial of funds')
  - any receiver contribution to usage through end to end clearinghouse
  - or bias fixed charges against receiving domain to compensate

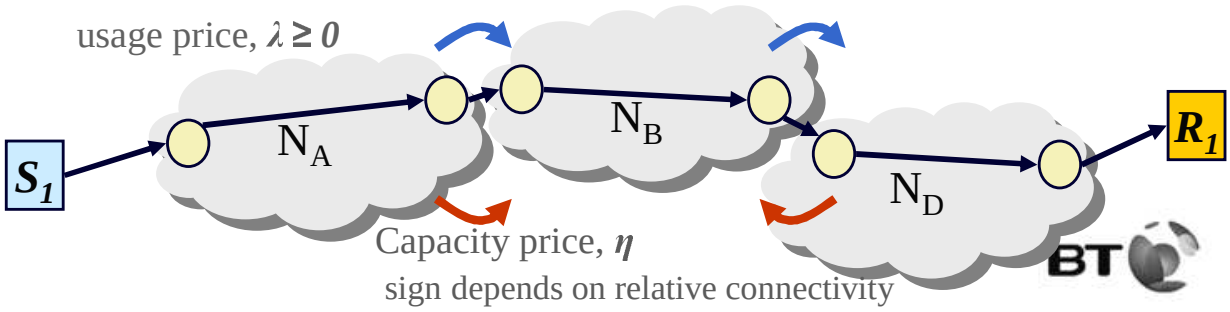

Q1. technical capabilities needed to support acceptable revenue models for providers?

first step: allow evolution of model

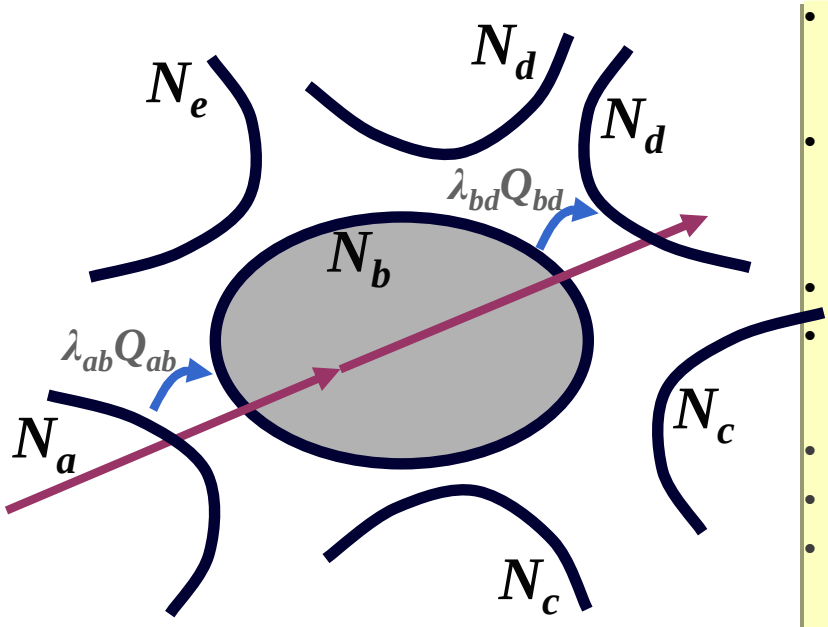

- decouple Q_{ab} from Q_{bd}
  - e.g Q_{ab} is volume
  - Q_{bd} is congestion
- common denominator is money
  - Profit attributable to flow, $\Pi_b = \lambda_{ab}Q_{ab} - \lambda_{bd}Q_{bd}$
- bulk pricing sufficient
- each price for rest of path from boundary to destination
- price effects localised
- contracts localised
- self-regulating, avoiding inter-carrier compensat'n (ICC) regulation
- global standards unnecessary

strong form: route agnostic
- `price for overall profit, win some, lose some
- or don't advertise loss-making routes

weak form: separate price for each subset of routes (e.g. all N_d)

BT

Q2. Constraints on pairwise agreements to support concatenated service?

minimum interconnect requirements (a)

- A2a) confine retail complexity to a higher layer e2e market
  - sender/receiver re-apportionment
  - roaming
- otherwise locks-in to single model for all interconnect
  - sufficient condition: interconnect contracts strictly bilateral (pairwise)

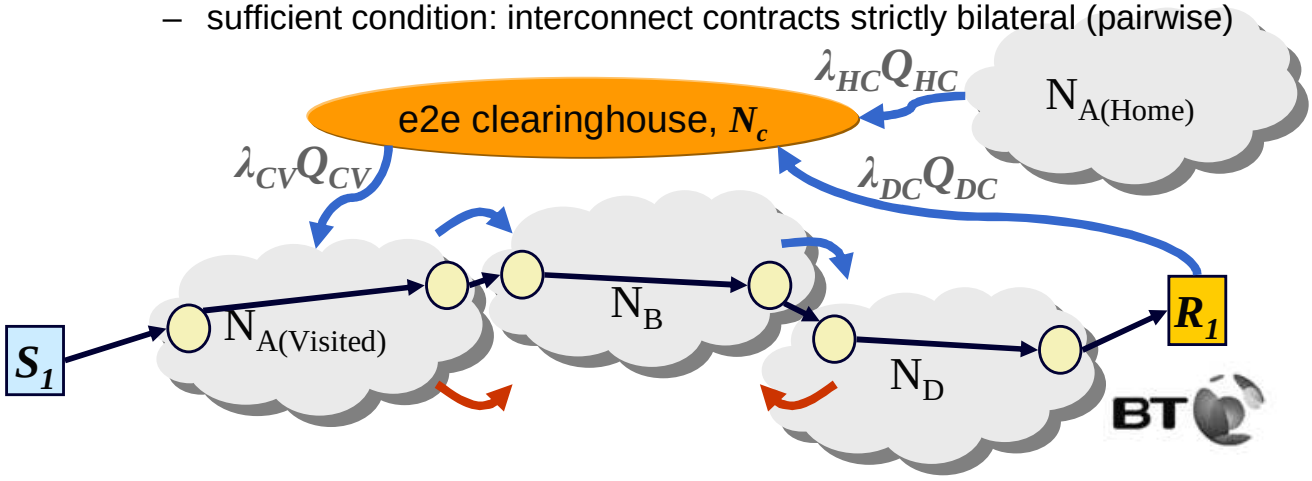

Q2. Constraints on pairwise agreements to support concatenated service?

minimum interconnect requirements (b)

- A2b) congestion pricing sufficient
  - can synthesise any QoS at edge, from congestion (ECN) pricing
  - simple, bulk, passive replacement for traffic policing
  - pushes back congestion upstream (cf. TCP)
- need longer slot to explain
  - simple, but unfamiliar territory for many
    - (cf 95th percentile peak demand or time of day volume pricing)
  - subject of IP QoS research since 1997
  - recently solved outstanding problems (to be proven)
    - direction of control (including routing/traffic engineering)
    - avoiding dynamic pricing in retail market

BT

interconnect QoS settlements – summary

- single model for end-game: congestion pricing

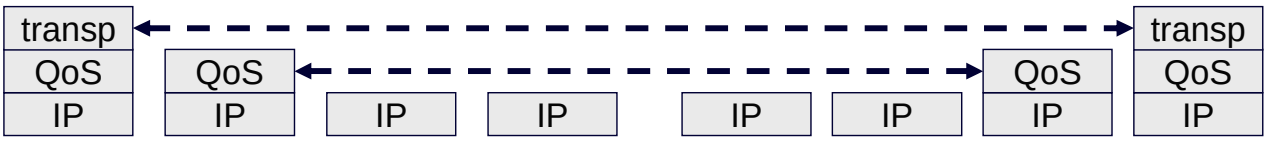

- or extra cost & revenue of more complex interconnect
  - to exploit temporary monopoly positions?

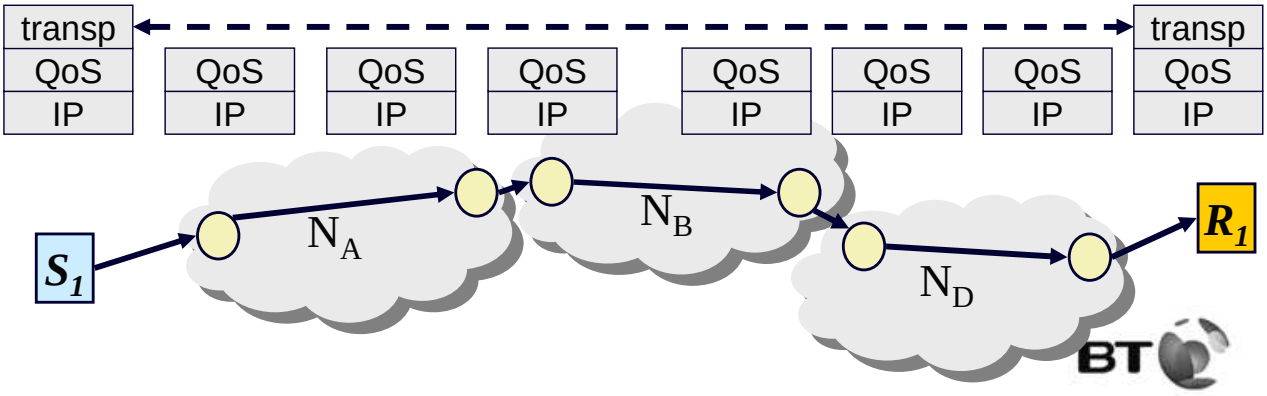

interconnect QoS – settlements
agreeing an industry model

- scope: the usage/QoS part of tariffs
- if we don't agree a layered industry model
  - it will cost us all hugely more to handle the mess
- alternatives within a single model:
  - only sender pays throughout network layer?
  - approx equal sender-receiver contribution throughout network layer?
- forum to agree this industry model?

BT

more info

- Bob.Briscoe@bt.com
- Paper
  - The Direction of Value Flow in Multi-service Connectionless Networks <http://www.m3i.org/papers/main.html#bt>

end-game: inter-domain congestion pricing

- **passive & extremely simple**
- recall sending domain pays to receiving domain $C = \eta X + \lambda Q$
- congestion charge, Q over accounting period, T_a is $Q = \Sigma^{T_a} \rho_i^+$
- ρ_i metered by **single bulk counter** on each interface
- **impairments trivial**