designing for tussle

case studies in
control over control

Bob Briscoe
BT Networks Research Centre
Jun 2004

role of communications research?

- pushing forward bounds of the possible
- help industry/society with comms technology choices
- to make an impact

- not just technical; also social, commercial
  - inseparable interwoven issues
  - ideal: multi-disciplinary expertise
  - sufficient: reasonable cross-discipline awareness
- otherwise will not make impact

BT

communications control

- problem: evolvability vs. infrastructure viability & abuse

end point control enables fast evolution of new scalable services 'e2e design principle'...
...but
- commoditises network operators who add value through service bundling
- requires end points to co-operate towards common goal

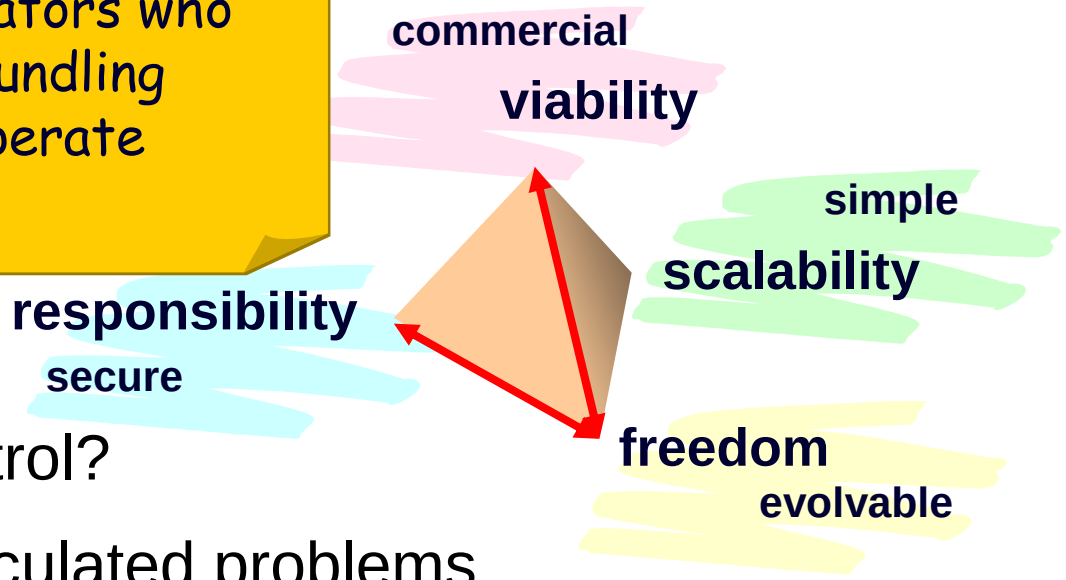

- who should be in control?
- DARPA NewArch articulated problems
  [BradenClarkShenkerWroclawski00]

BT

control assumptions: examples

- authentication: who checks id?
- denial of service attack or congestion?: who decides?
- resource sharing: who decides fairness criterion?
- peer to peer sharing/ad hoc: why share resources?
- end-point vs. middle control: purely technical?

- aim to explicitly state control assumptions

BT

control assumptions in typical papers

- neutral ← not so common
- unformed ← fine
- unconscious ← worst
- conscious unstated ← rarely succeeds
- conscious stated ← fine

- control over control ← subject of this talk
  - decide control model at run-time, not design time
  - improve infrastructure evolvability and viability...

BT

evolution of evolvability research

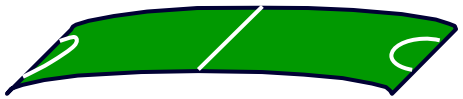

❶ end to end arguments [SaltzerReedClark84]

- protect generic investment, surrender control to foster innovation

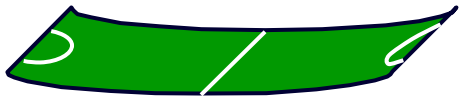

❷ end of e2e [ClarkBlumenthal00]

- ends not trusted to co-operate with whole
- middle needs investment incentive

❸ end of (end of e2e) [Shenker, Kelly, Varian, Crowcroft, Anderson etc]

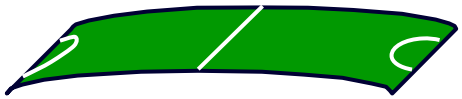

- game theoretic mechanism design

❹ argument is the end [ClarkSollinsWroclawskiBraden02]

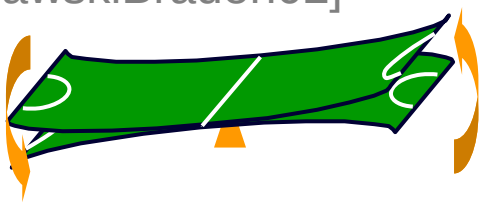

- design for tussle

BT

comms infrastructure control
a history of tussle

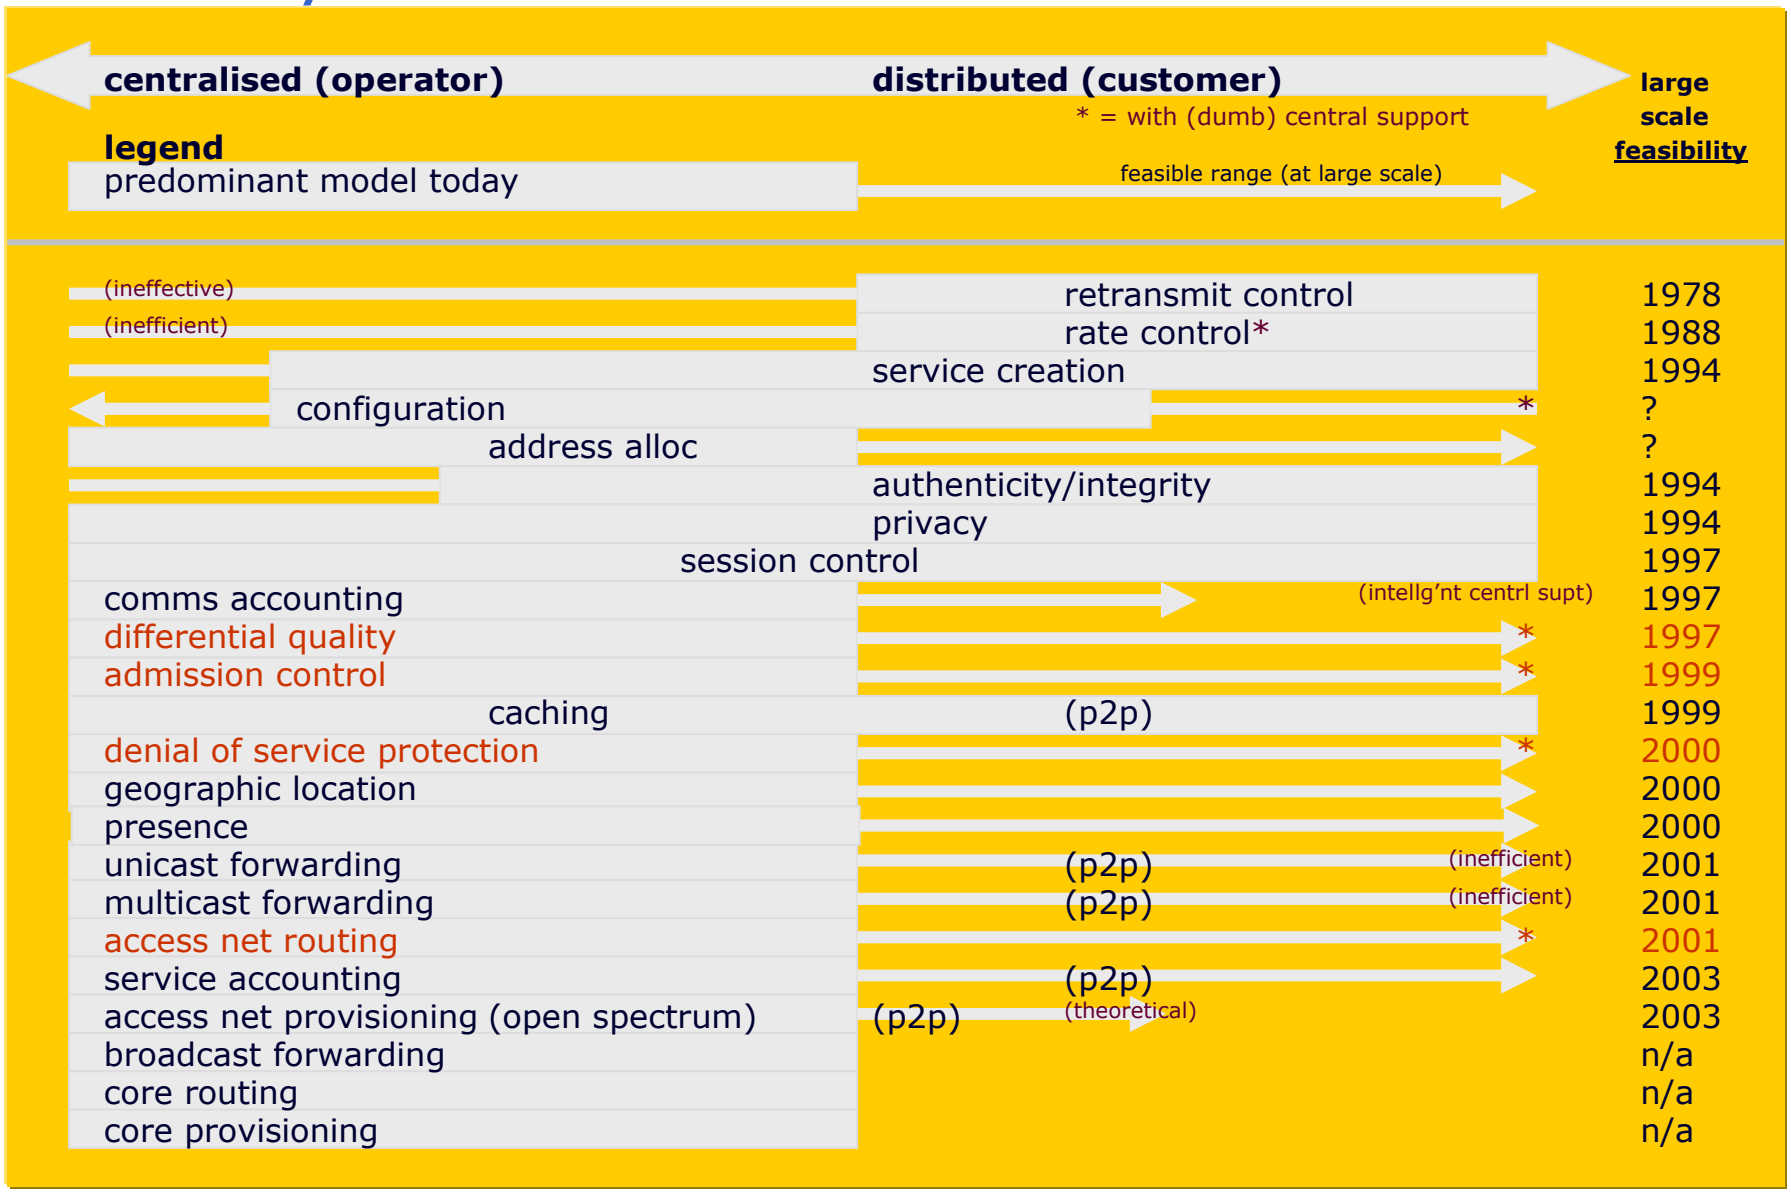

spectrum of control

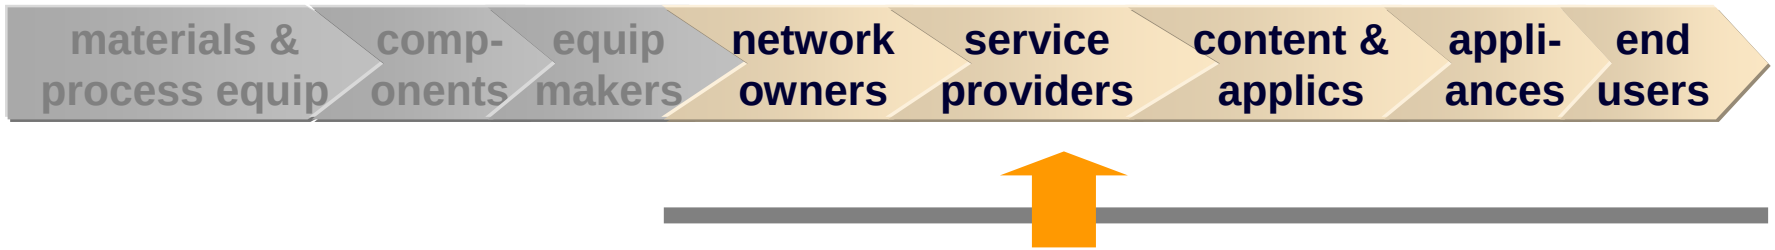

- having designed for extremes
- can also move control to intermediate points

case study: denial of service mitigation

❶ e2e: iTrace: ends: detection& trace; middle: previous hop

- 1:1M data packets trigger ICMP iTrace packet at each router
- message to dest address giving present & previous hop address
- dest under attack can trace back to earliest honest address on path
- push-back filters into network

❷ e2e problems

- ends not trusted: spoof attack to install false filters
- middle needs incentive to invest in iTrace upgrades

❸ e2e fixed

- authenticate filter requests hop by hop

❹ design for tussle

- move detection & trace to proxy one notch in from ends

case study: denial of service mitigation

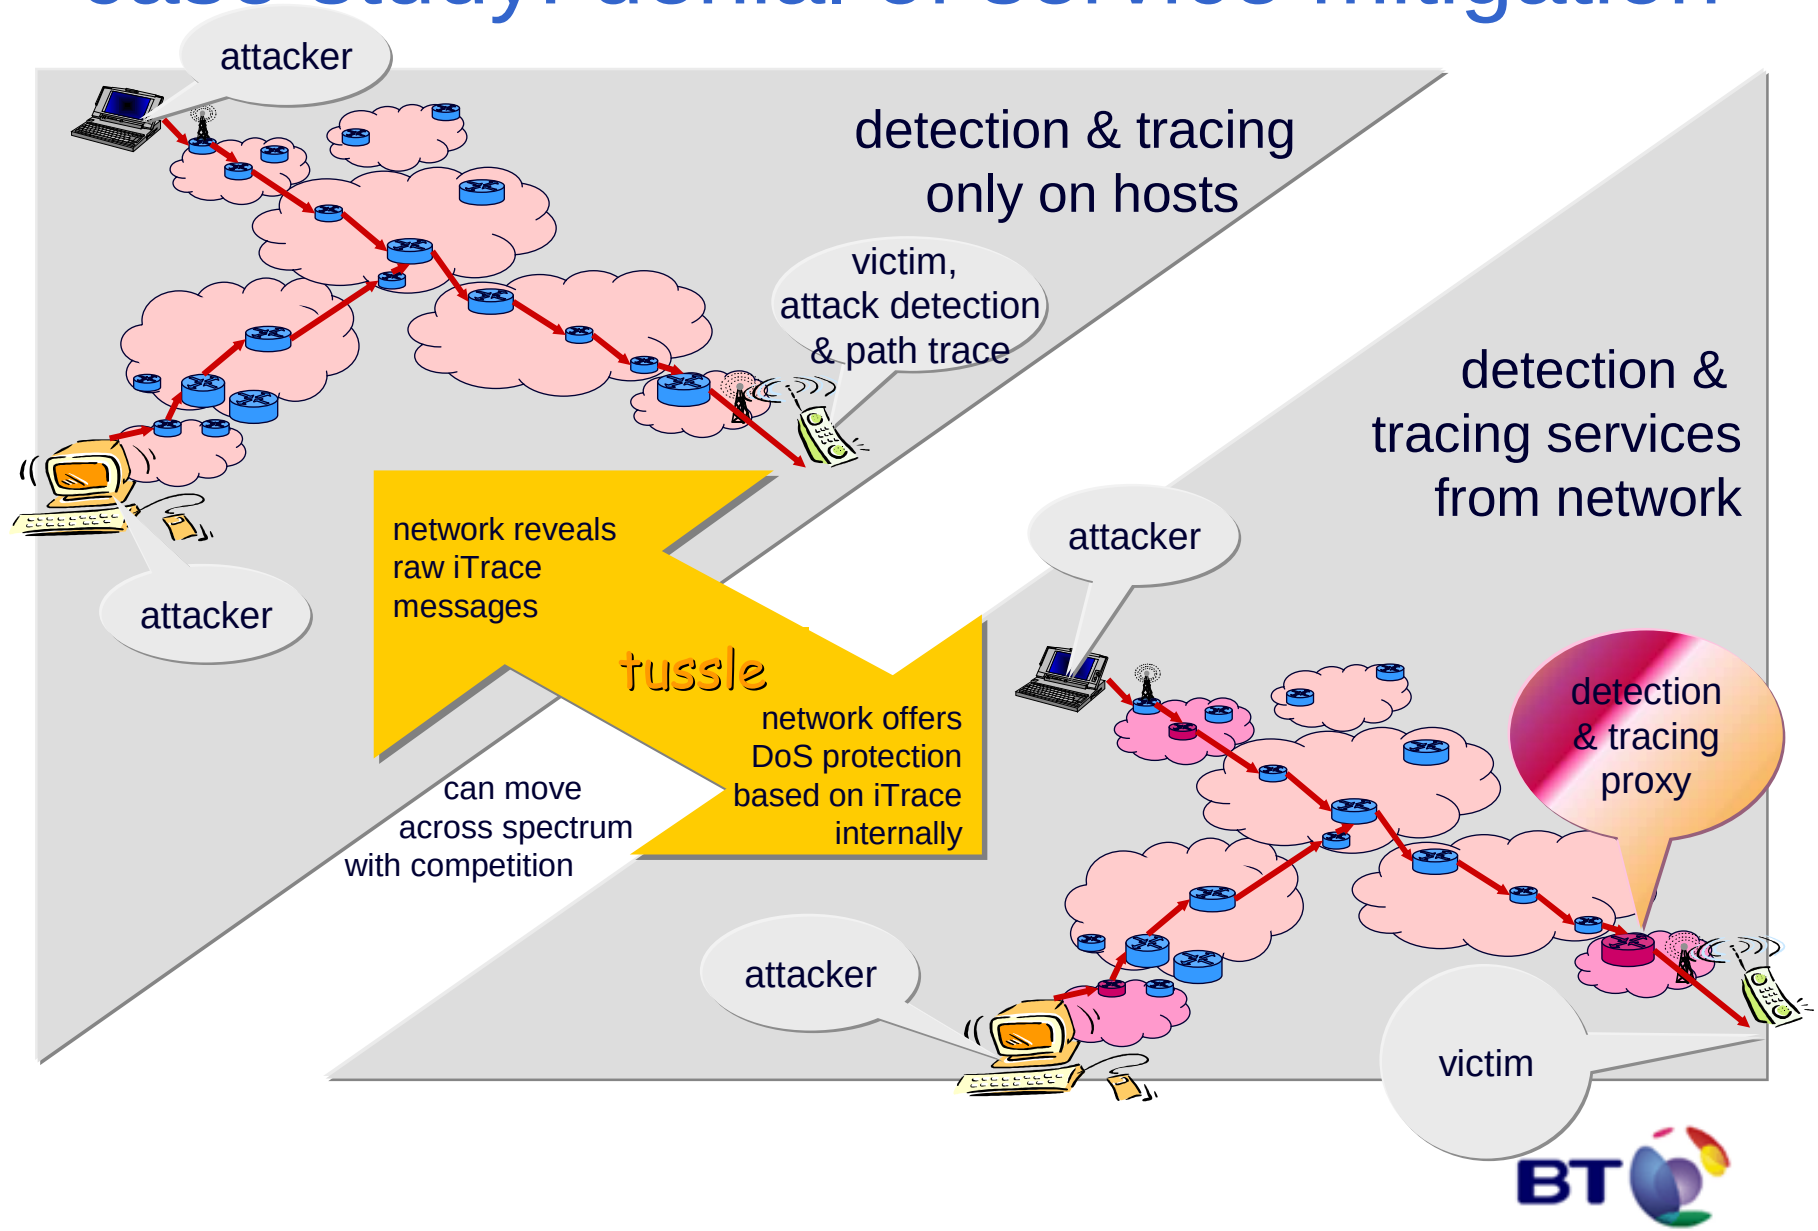

case study: contractual mobility

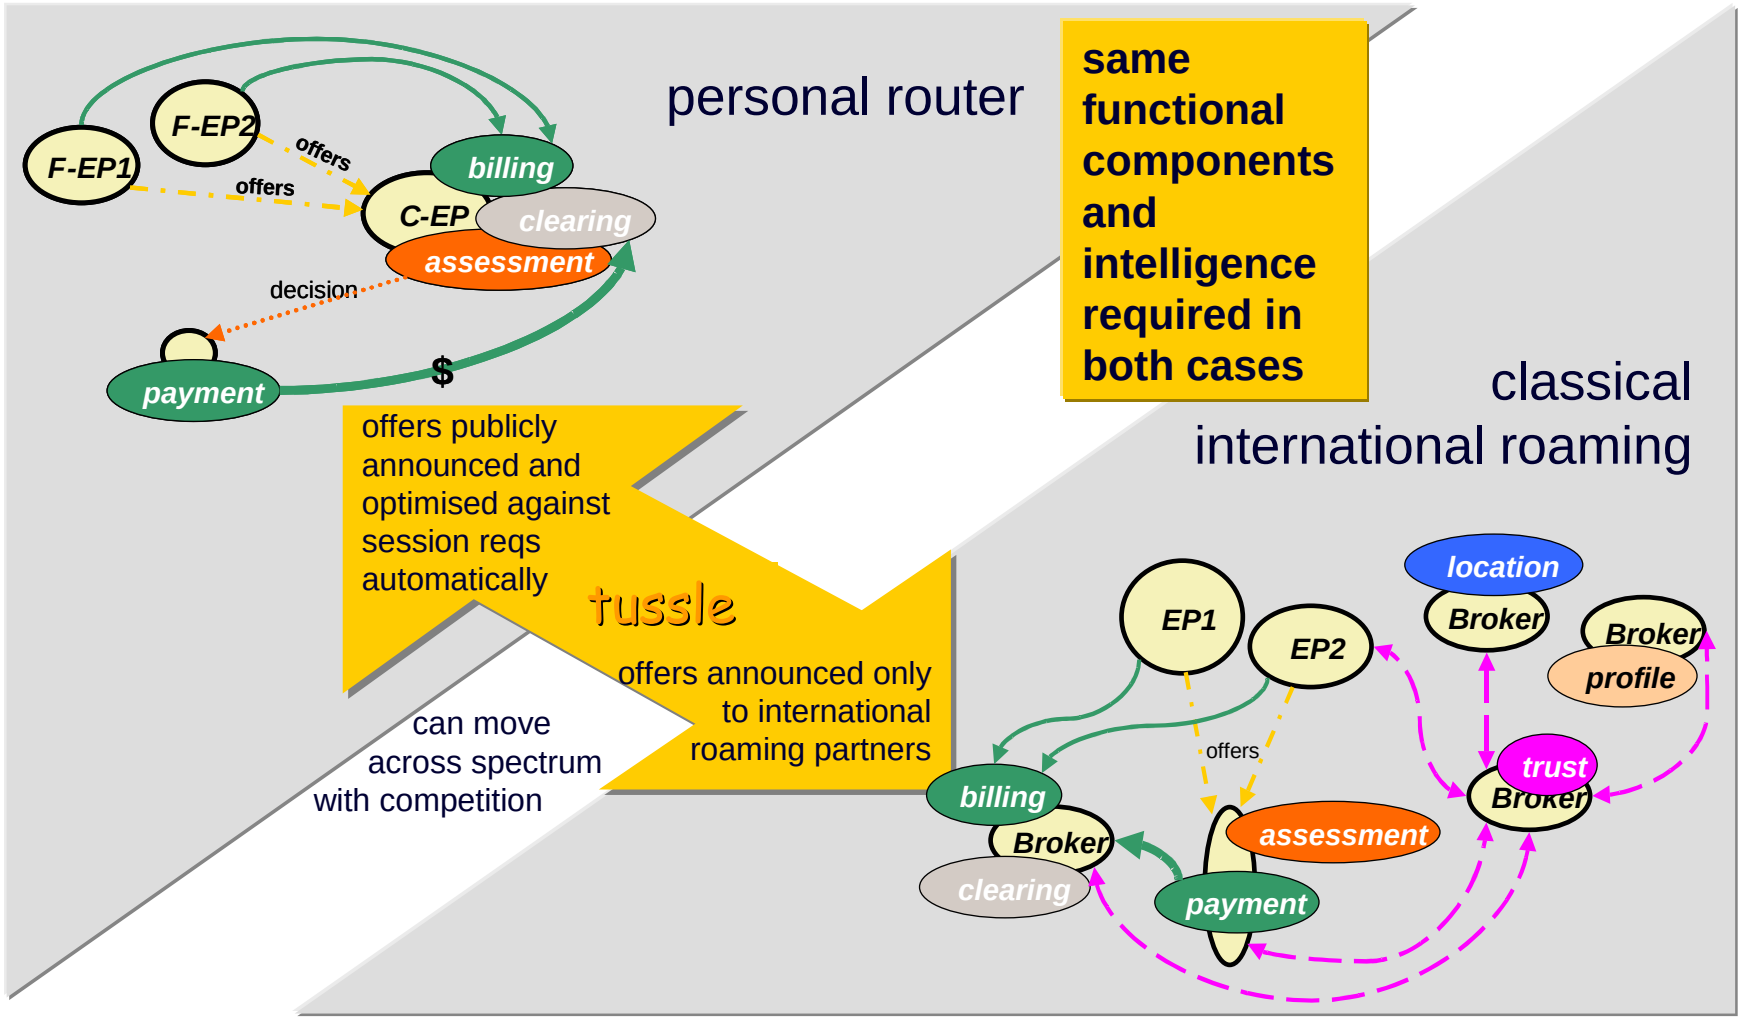

case study: quality of service

❶ e2e: TCP/IP: ends: congestion control; middle: forwarding

- transmission control protocol (TCP) [VanJacobsen88]
  explicit congestion notification (ECN) [Floyd94]

❷ e2e problems

- ends not trusted: VoIP free-riding
- middle needs investment incentive
  Intserv [BradenClarkShenker94], Diffserv [ClarkWroclawski97]

❸ e2e fixed

- shadow pricing, proportional fairness [GibbensKelly99]

❹ design for tussle

- guaranteed QoS synthesis [Karsten02]
- control over control [Briscoe02]

BT

QoS context: cost realities

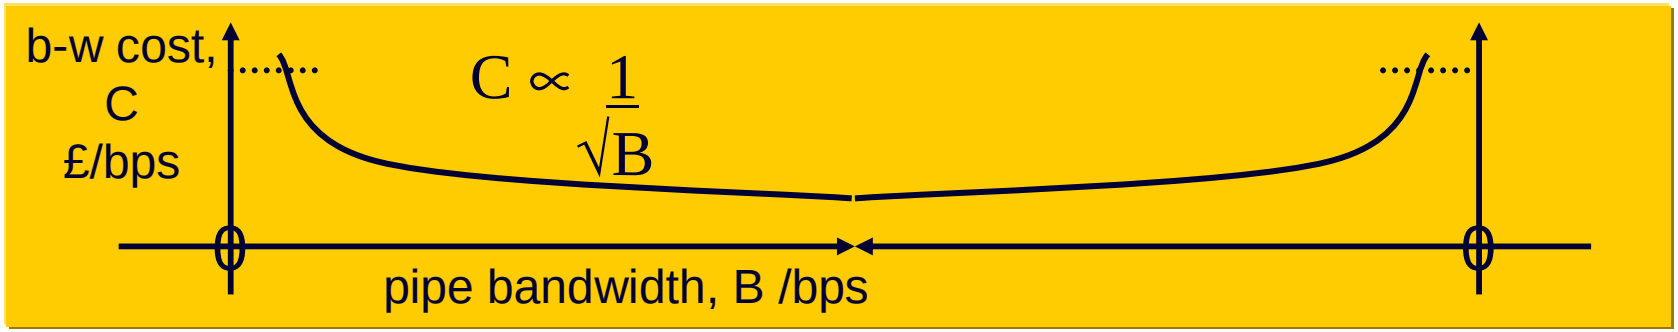

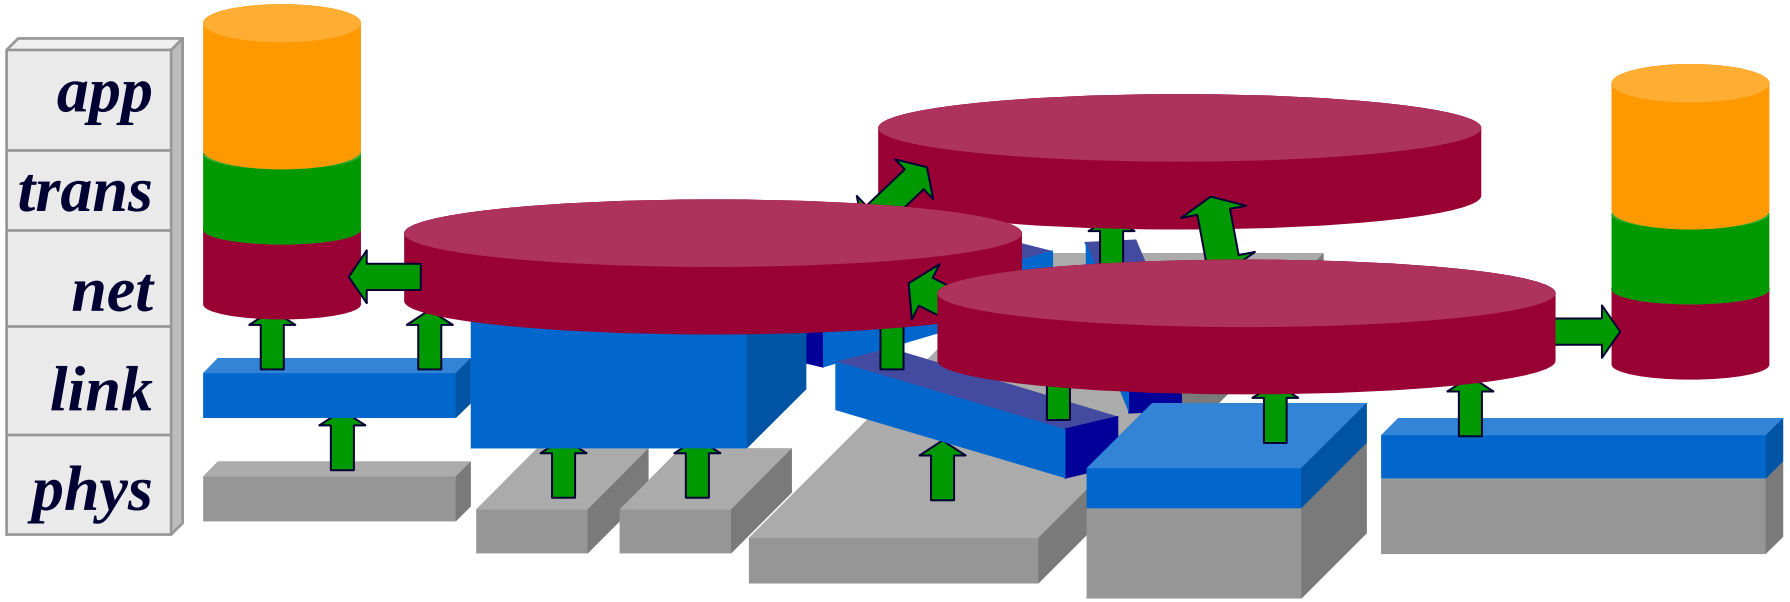

❶ e2e design

TCP: business model

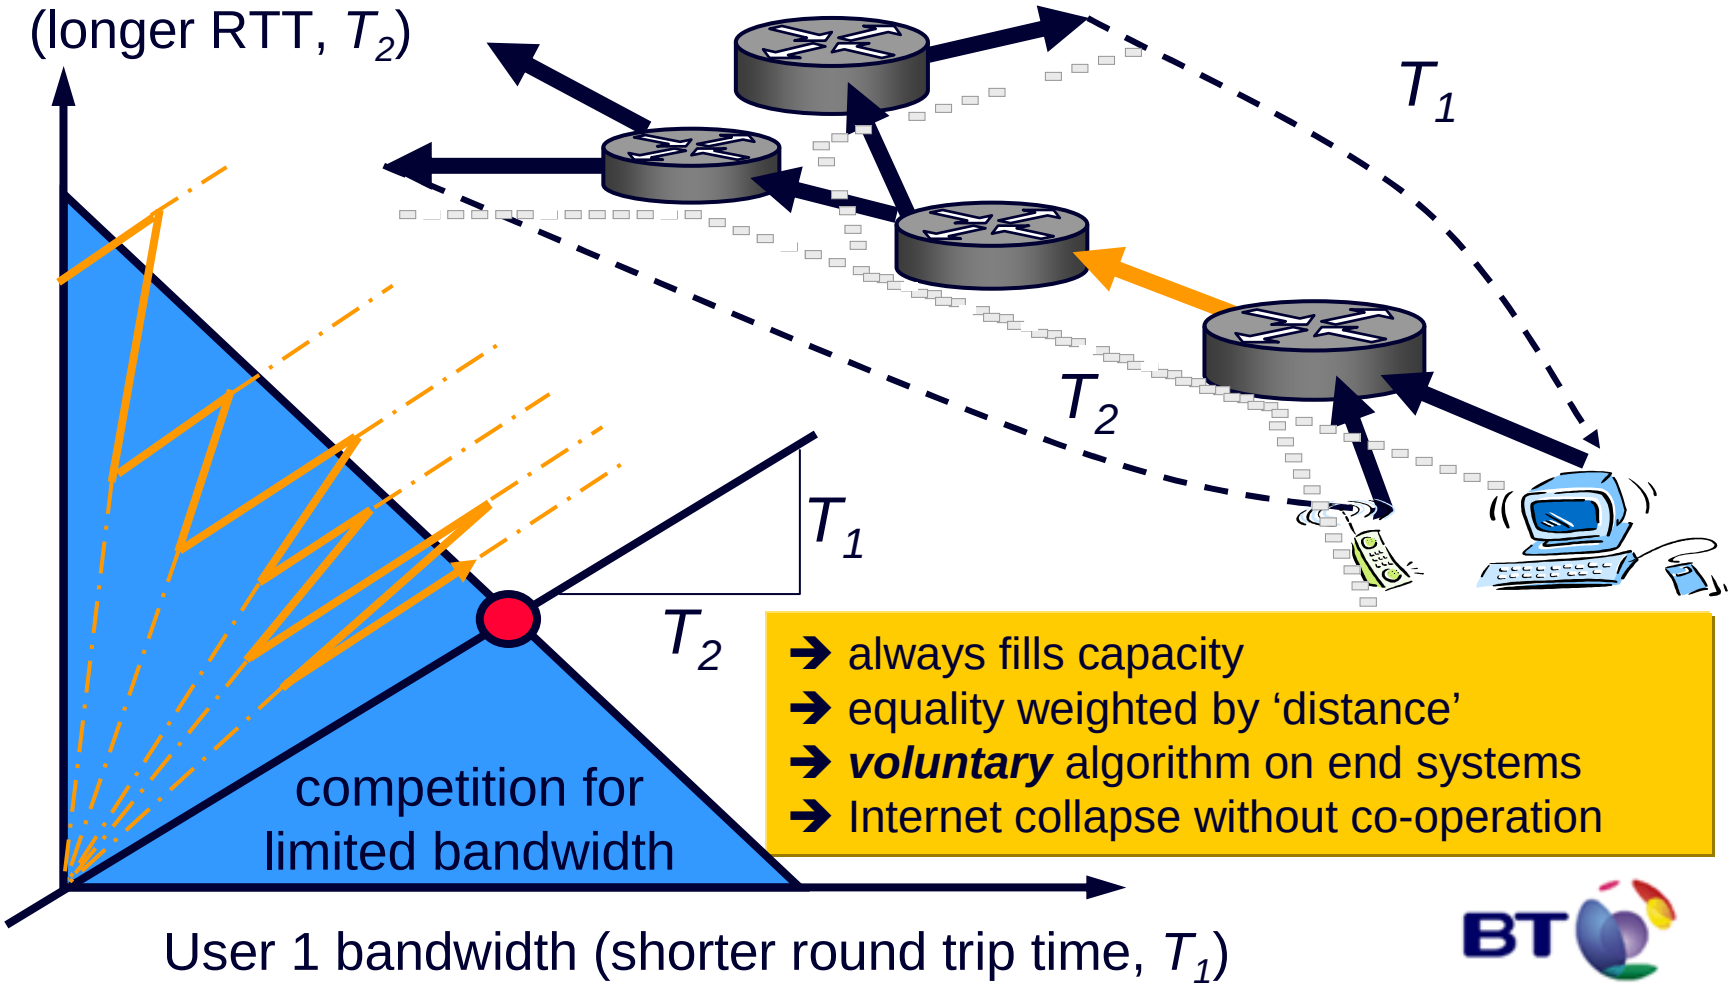

- ➔ always fills capacity
- ➔ equality weighted by ‘distance’
- ➔ ***voluntary*** algorithm on end systems
- ➔ Internet collapse without co-operation

BT

❷ e2e problems

greed breeds policing

- voice over IP
  - if experience congestion, send more
- integrated services
  - users reserve path resources (ReSerVation Protocol)
  - networks control admission then police traffic
- differentiated services
  - provision prioritised logical classes of service
  - traffic classified (Diffserv field in IP) and policed
  - congestion avoided for higher classes, usually
- middle takes control
  - can vertically integrate with media business

❸ e2e gets fixed

explicit congestion notification (ECN)

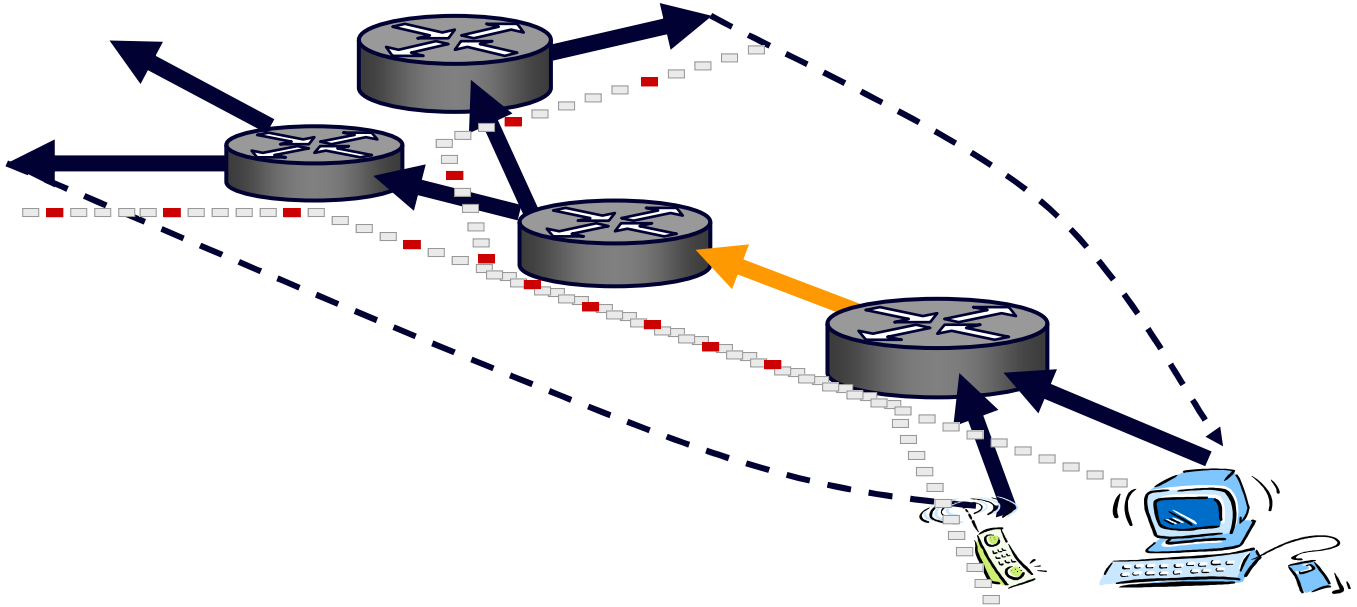

- **without ECN: first sign of congestion is loss**
- **with ECN: mark packets randomly as congestion builds**

- **2001: ECN standardised into IP & TCP**
- extensible for marking *before* congestion onset (virtual queue)

BT

❸ e2e gets fixed

DIY QoS

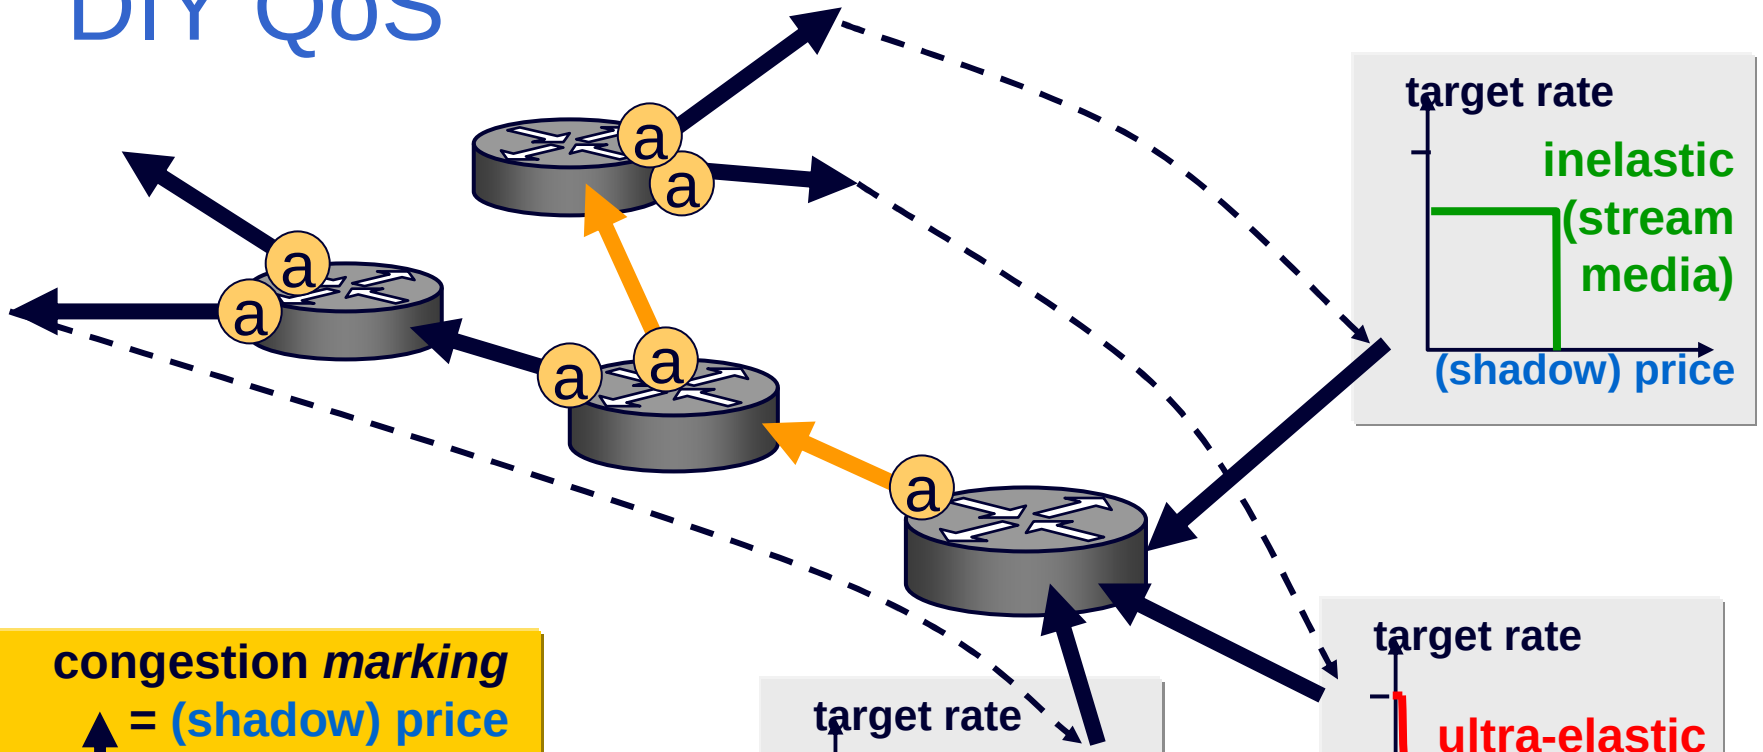

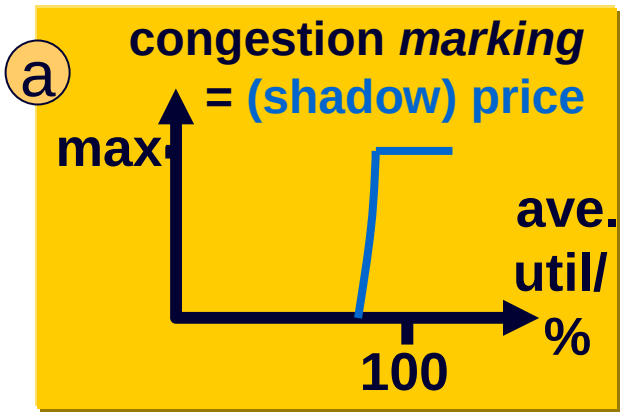

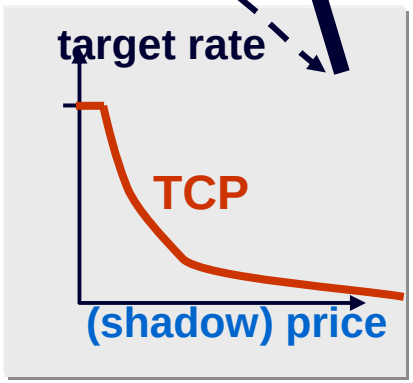

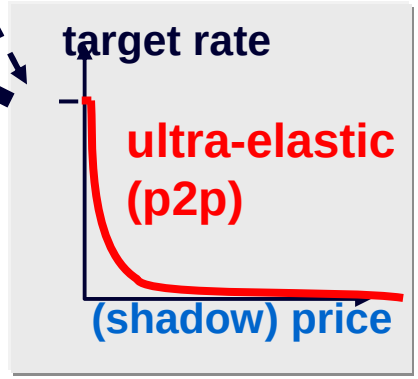

BT

❹ design for tussle

guaranteed QoS synthesis

IP routers	Data path processing
Reservation enabled	① Reserved flow processing
RSVP/ECN gateway	② Policing flow entry to CP ④ Meter congestion per peer
ECN only	③ Bulk ECN marking CP prioritised over BE

- guarantees over simple middle
- allows vertically integrated media business at edge
- DIY QoS one notch in
- uses 3 QoS standards but not their architectures
- PSTN replacement but evolvable business model...

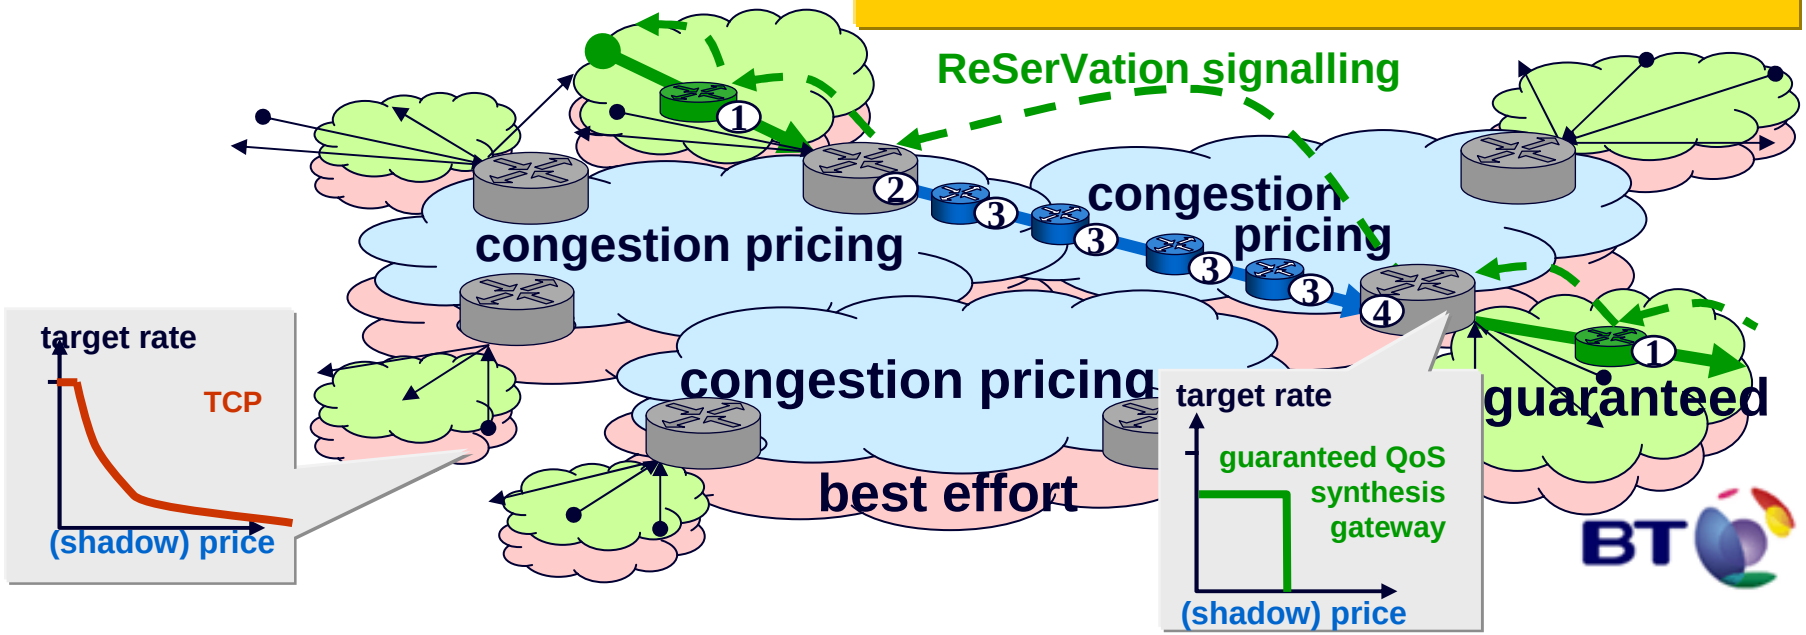

case study: QoS

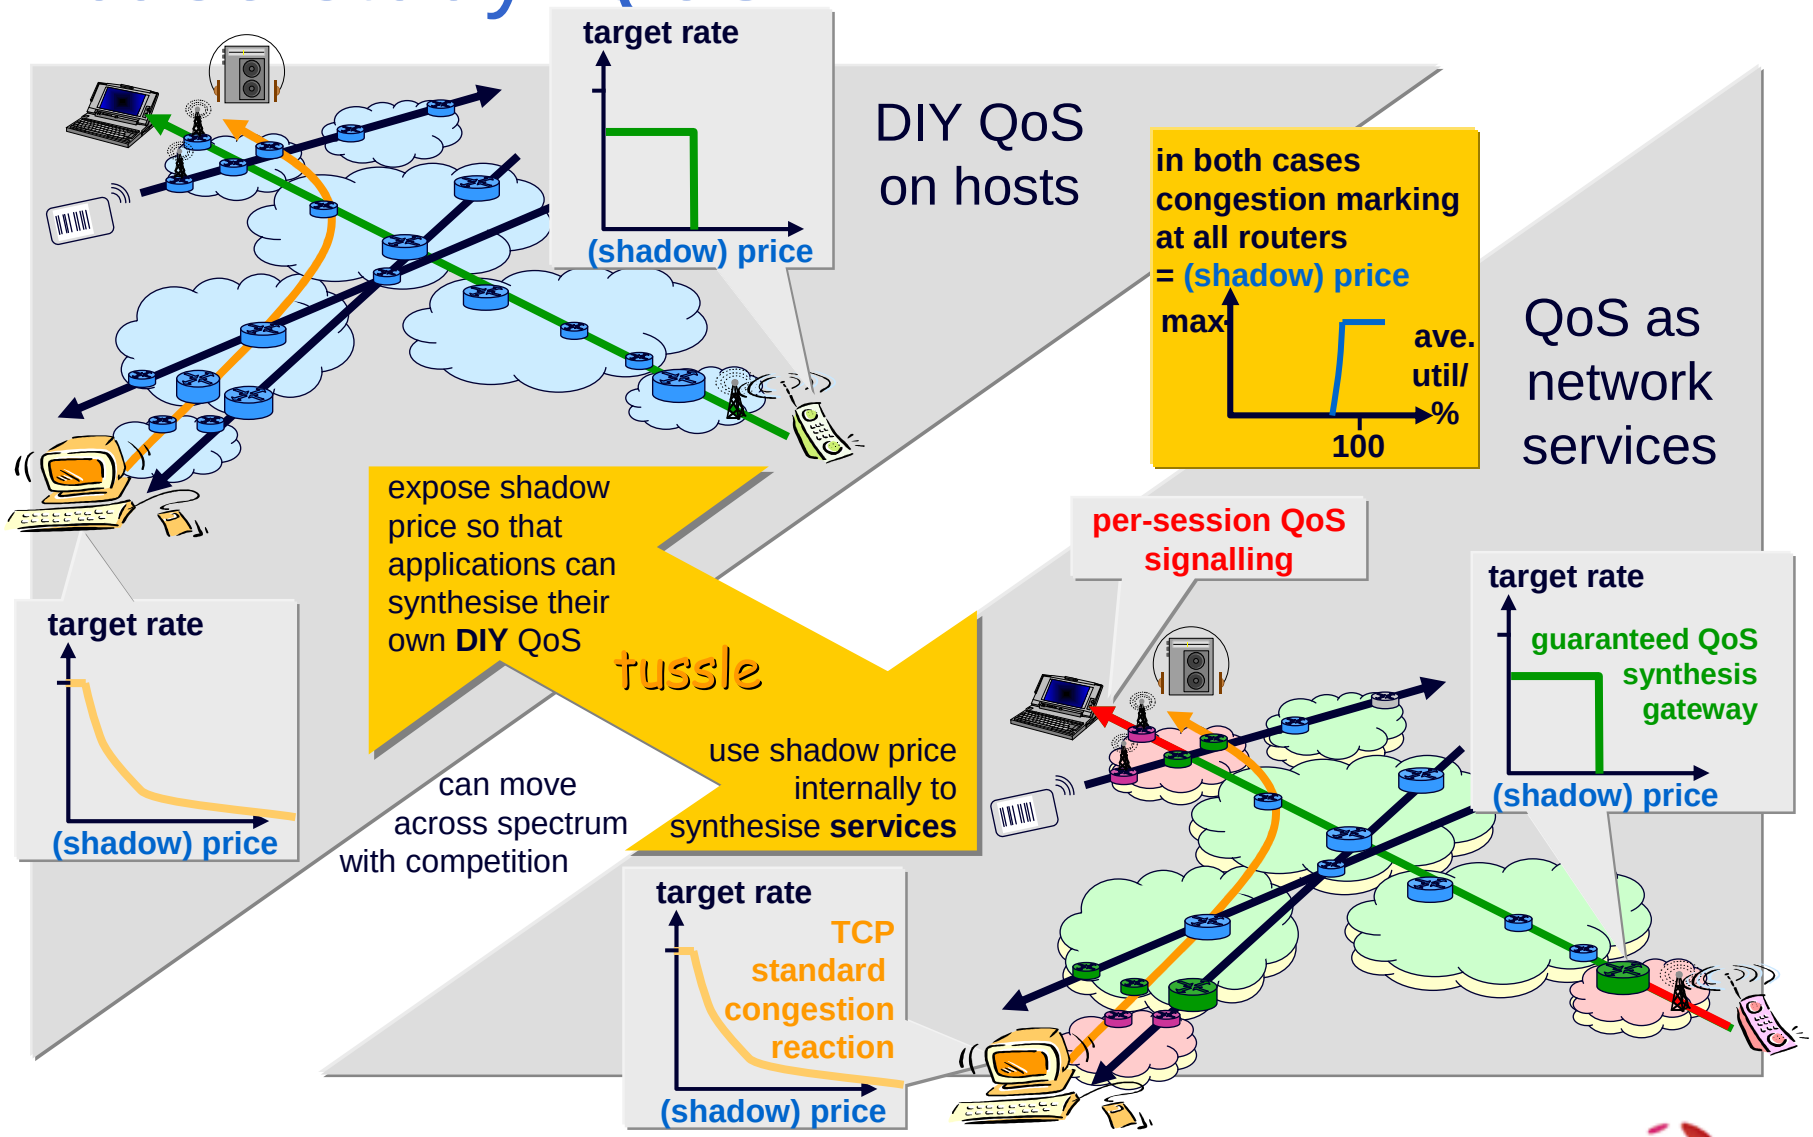

case study: multipoint messaging

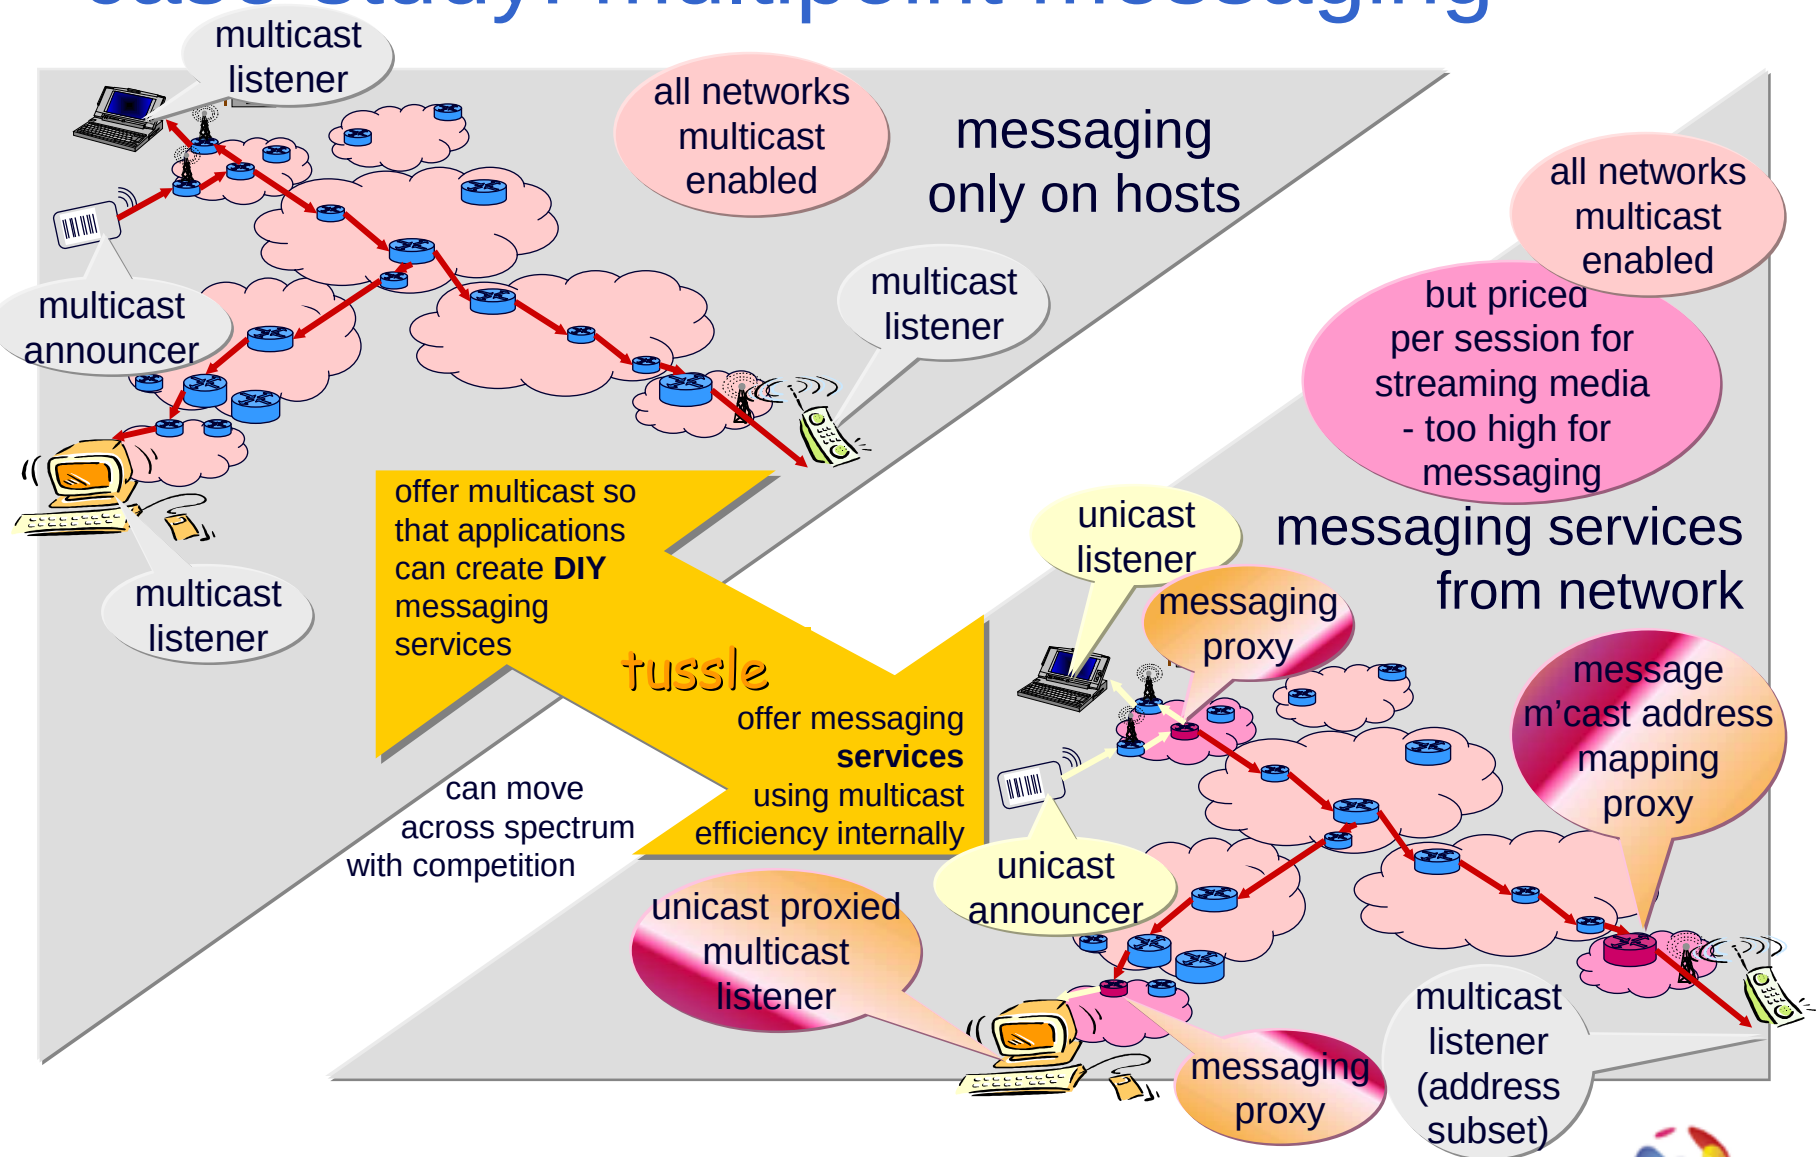

lesson: downstream knowledge upstream

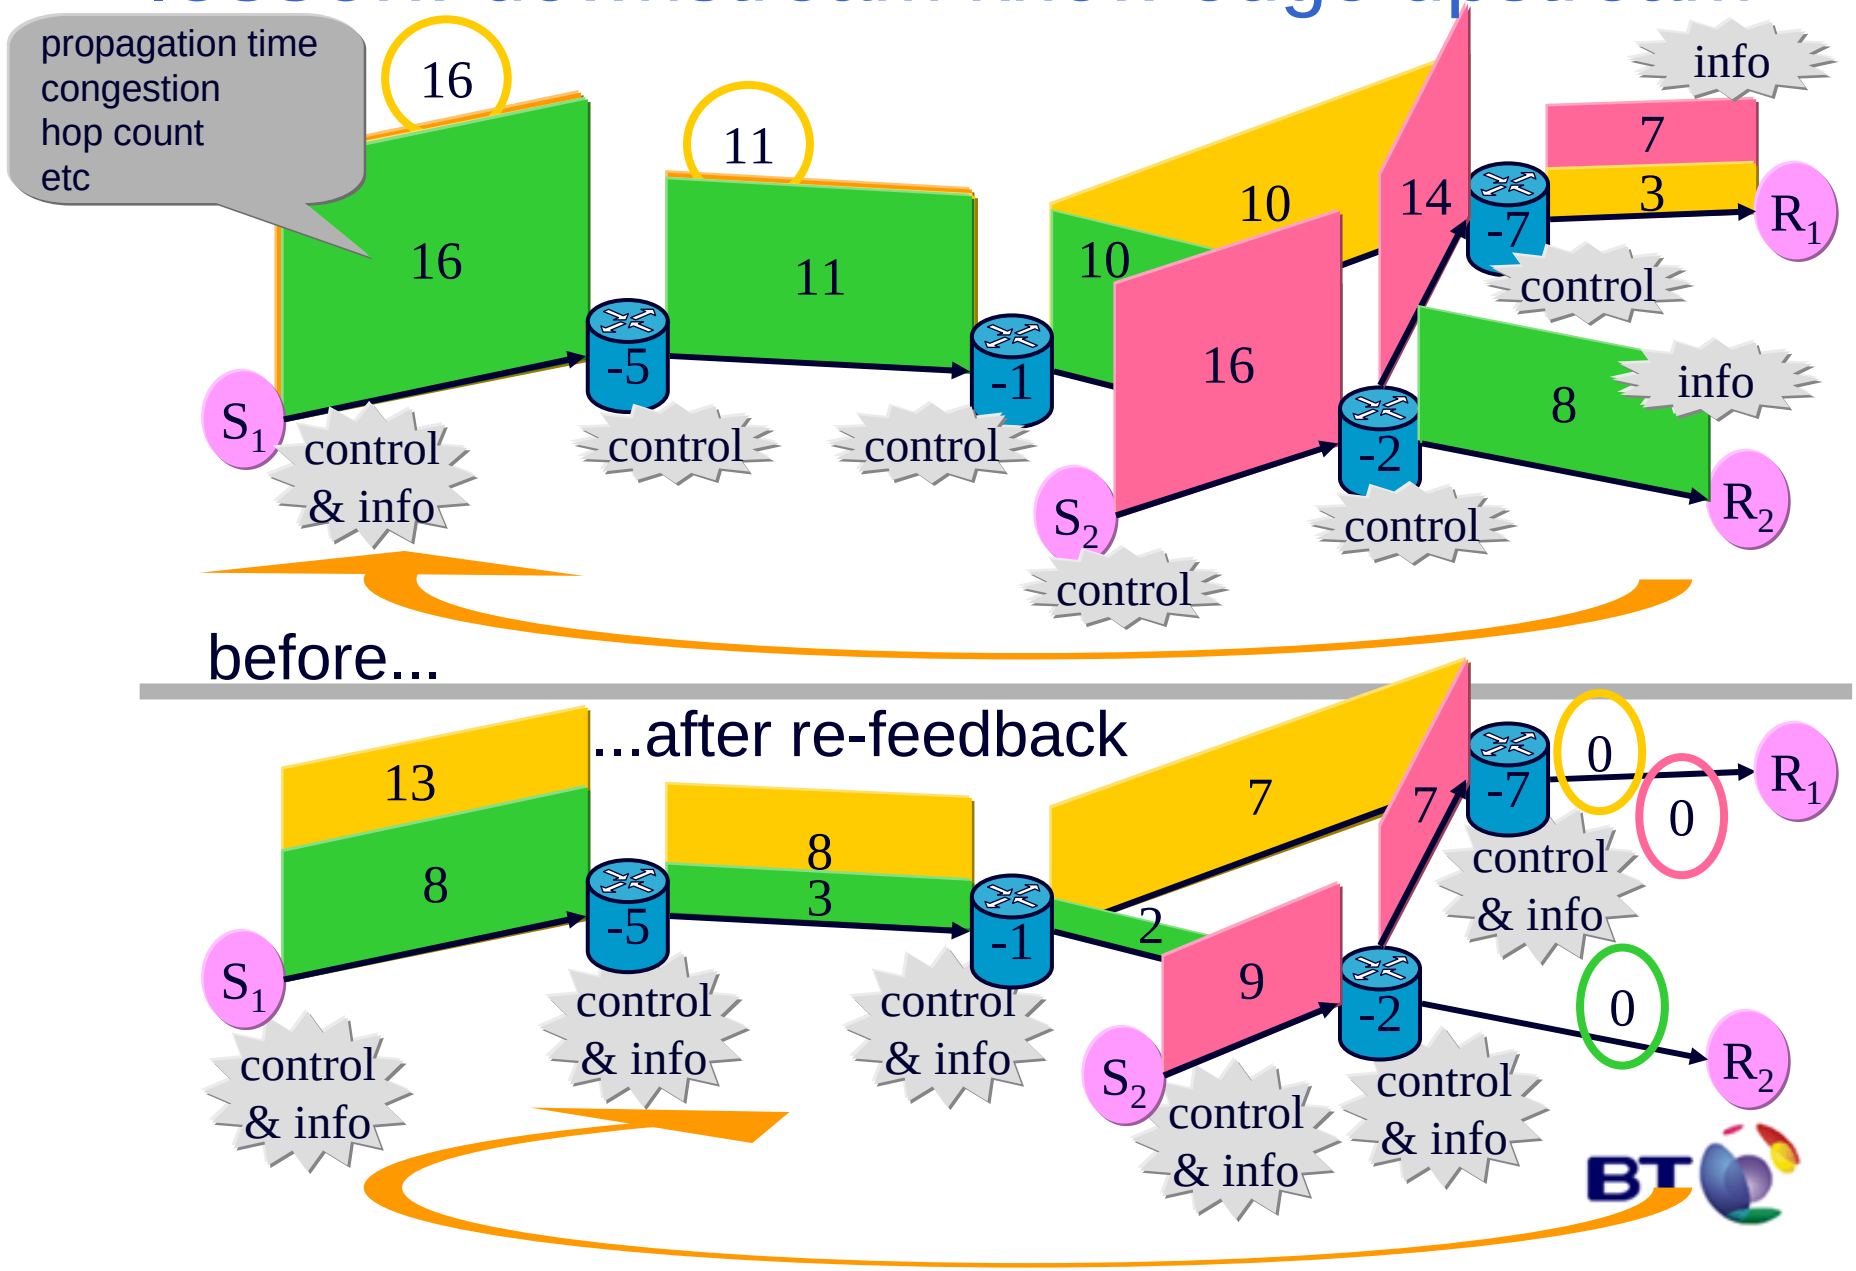

packet re-feedback

downstream knowledge upstream

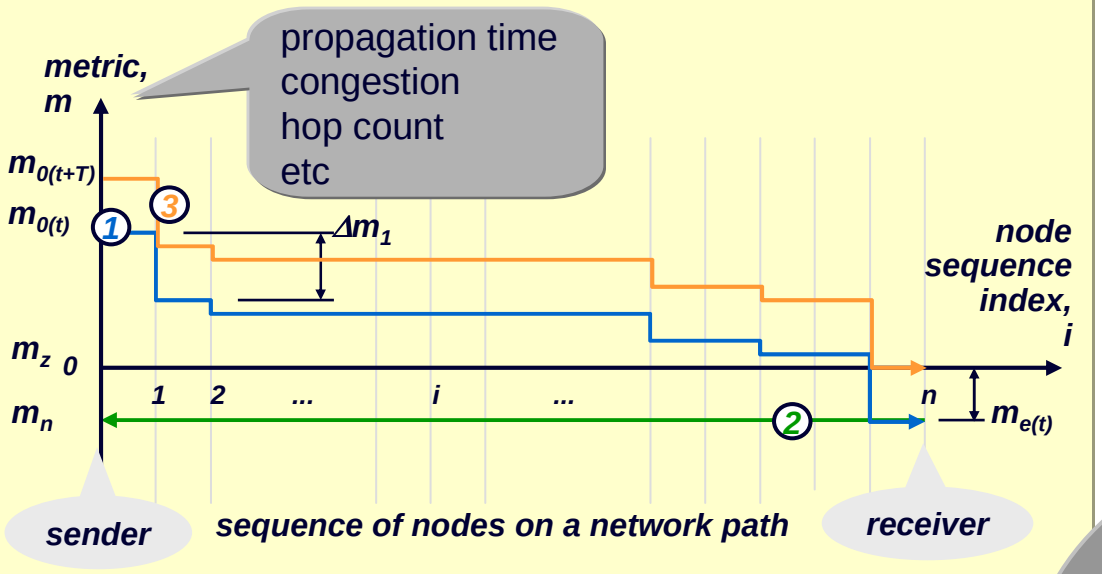

all network nodes know as much about downstream path as data source
- level playing field

incentives to hit target

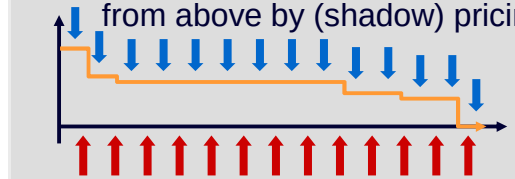

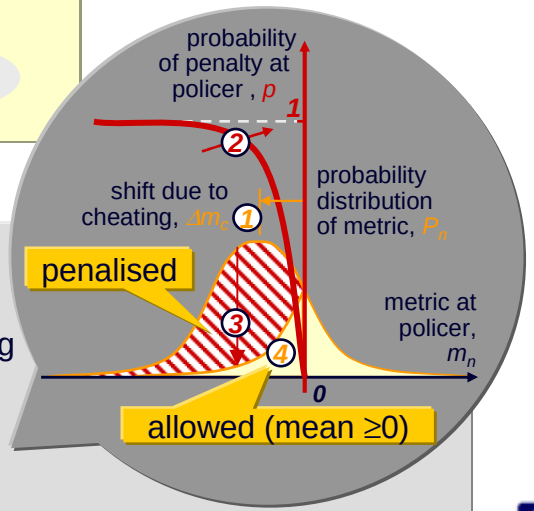

control over control?

- control can migrate 

- sell different control models to different markets
  - DIY and “do it for you” customers
- 

can re-sell control package each time

- how to control where control is?
  - offering protocol response at a price ‘switches on’ its importance
- what controls where the control is?
  - market advantage, competition
  - regulation

summary of approach

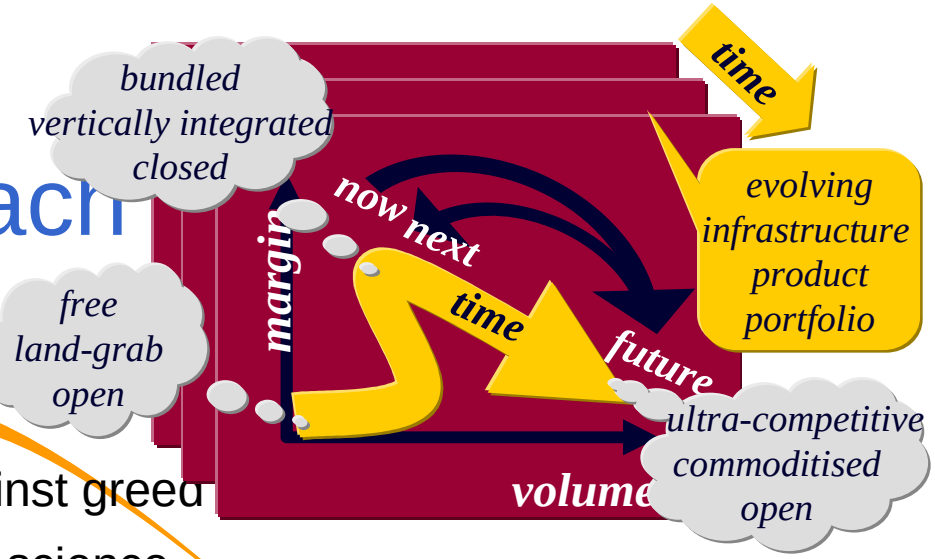

- design *as if* e2e
  - include proofing against greed
  - based on underlying science
- design edge interception of e2e protocols
- let the tussle commence
  - capture market share with free, open product
  - pull in control from ends to edge
  - competition gradually commoditises
  - giving up control stimulates new innovation
- layer under next product

Net-head heart
Bell-head skins

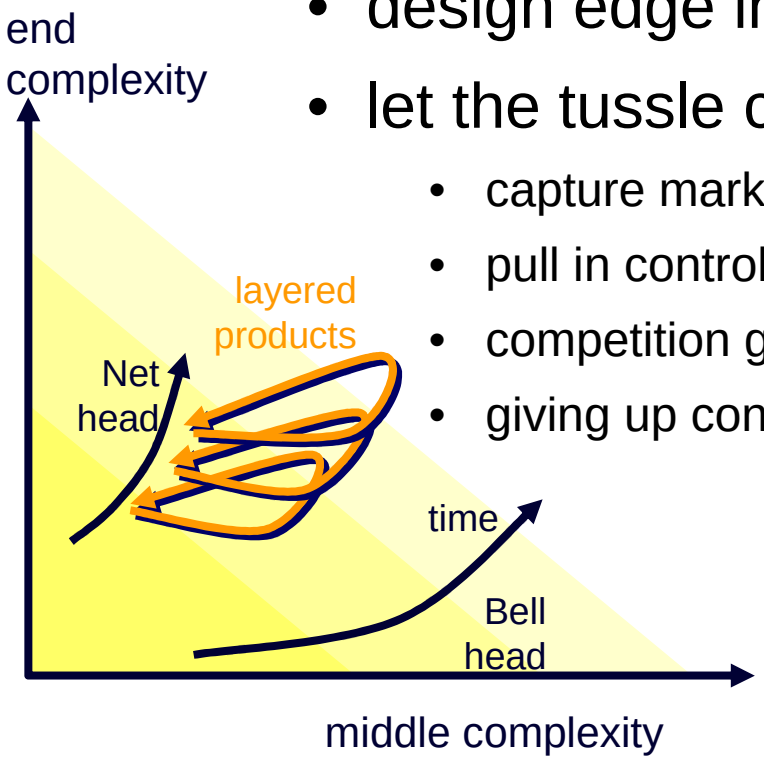

research agenda implications

- pure technical research sometimes valid
- but often implicit commercial assumptions missed
- encourage articulation of commercial assumptions
- encourage multi-disciplinary research
  - at fundamental level, not just applications

questions?

control over control

further info

- Bob.Briscoe@bt.com

- [SaltzerReedClark84] Jerome H. Saltzer, David P. Reed, and David D. Clark, “End-to-end arguments in system design,” ACM Transactions on Computer Systems, 2(4):277–288 (Nov 1984)
- [GibbensKelly99] Richard J. Gibbens and Frank P. Kelly. Resource pricing and the evolution of congestion control. Automatica, 35, URL: http://www.statslab.cam.ac.uk/~frank/evol.html (1999)
- [ClarkBlumenthal00] David Clark and Marjory Blumenthal, “Rethinking the design of the Internet: The end-to-end arguments vs. the brave new world,” In Proc. Telecommunications Policy Research Conference (TPRC’00), URL: http://www.tprc.org/abstracts00/rethinking.pdf (Sep 2000)
- [BradenClarkShenkerWroclawski00] Bob Braden, David Clark, Scott Shenker and John Wroclawski, “Developing a Next-Generation Internet Architecture,” DARPA White paper, URL: http://www.isi.edu/newarch/DOCUMENTS/WhitePaper.pdf (Jul 2000)
- [Briscoe02] Bob Briscoe, "M3I Architecture PtI: Principles” Deliverable 2 PtI, M3I Eu Vth Framework Project IST-1999-11429, URL: http://www.m3i.org/results/m3idel02_1.pdf (Feb 2002)
- [ClarkSollinsWroclawskiBraden02] David Clark, Karen Sollins, John Wroclawski and Robert Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet,” In: Proc. ACM SIGCOMM'02, Computer Communication Review 32 (4) URL: http://www.acm.org/sigcomm/sigcomm2002/papers/tussle.pdf (Aug 2002)

discussion

- design for tussle is subtle
  - takes years of hindsight to get right
  - too late for early market advantage?
  - open, free land grab gives some breathing space
  - can tendering process cope with subtlety?
- does designing for commoditisation bring it forward?
  - is having no plan B more risky?
- parallels in Microsoft product evolution?
  - BIOS, DOS, Win, COM, .NET, Office

spare slides

Bob Briscoe

seamless resource control

😐 **traditional (optional):**
optimise ea subnet separately
e.g. Diffserv (open-loop)

☺ **new (required):**
optimise all paths together

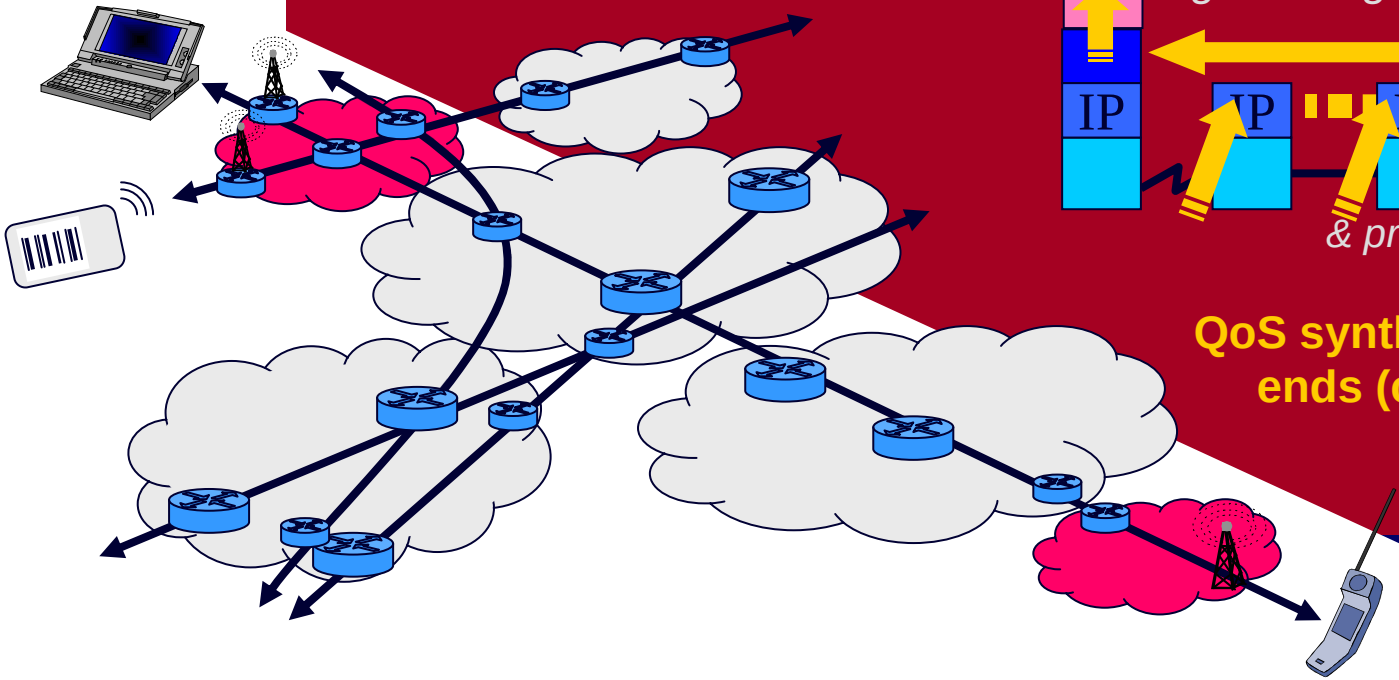

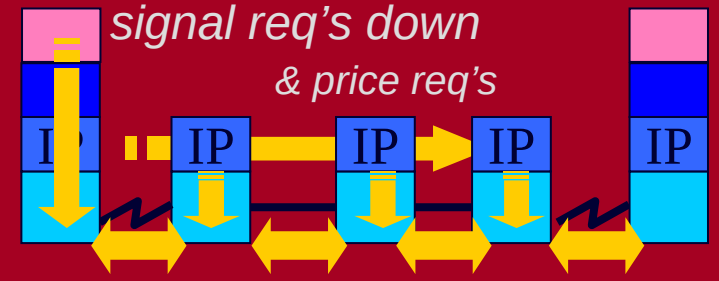

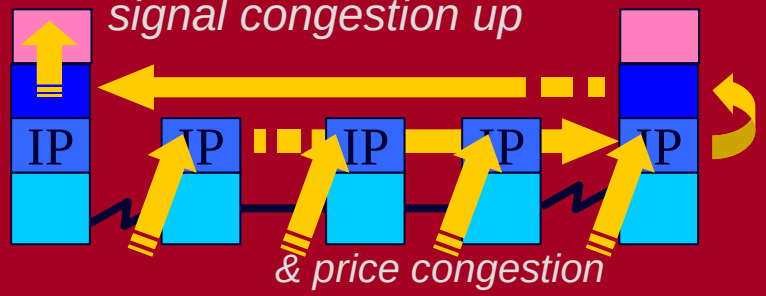

QoS synthesised by the ends (closed-loop)

Internet (not telco) industry approach

- creating *x*-like systems out of un-*x*-like parts
  - where *x* is some desirable attribute
- creating secure systems out of insecure parts
- creating reliable systems out of unreliable parts
- creating intelligent systems out of unintelligent parts
  - eg. intelligent session control without an intelligent network
- creating QoS control systems out of non-QoS controllable parts
- creating a telephony system out of best effort Internet parts
- ...
- creates low cost systems out of low cost parts
- but the approach puts all the smarts at the ends, which...
- creates ~~profitable~~ broken value chains out of unprofitable players...?

BT