

Why are all architectural problems  
from 2000 still unsolved?



How would we know we had solved  
socio-economic problems anyway?

Bob Briscoe  
Chief Researcher, BT Group  
Apr 2007



# you can't have your dessert until you've eaten your vegetables

- careful not to invent problems to fit the research we want to do
- research agenda since DARPA NewArch (2000) all still unsolved
  - 'solved' = rough consensus and deployable code (ideally all solutions coherent)
- routing, naming, addressing (n)
  - policy controls on inter-provider routing
  - robustness & availability, inc mobility
  - reachability through middleboxes
- management (0)
  - policy-driven auto-configuration
  - failure management
- resource control (0)\*
  - highly time-variable resources
  - capacity allocation
  - extremely long propagation delays
- security (n)
  - attack resilience
  - traceability
- heterogeneity – cross-cutting agenda
  - enabling conflicting socio-economic outcomes (0)
  - enabling a variety of technical outcomes (n)

2

\* resource control:  
0 projects in NSF NeTS FIND  
1 retrospective paper in SIGCOMM'06

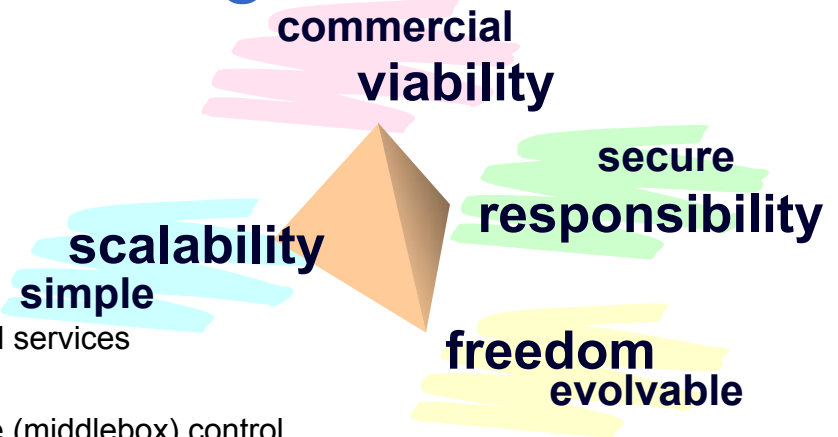


# networks research – enduring tensions

## design for tussle

between outcomes in this space

- not just self-supply (p2p, ad hoc)
  - but co-existence of ad hoc and managed services
- not just endpoint control
  - but co-existence of end control and edge (middlebox) control
- not just individual security / privacy
  - but co-existence of individual freedom and social/corporate control
- balance between approaches determined by natural selection
  - market or social (e.g. government) control
  - society & the economy: shaping the Internet and shaped by the Internet
  - requires multidisciplinary research teams
- imposing your political values through your design
  - just means your design will get distorted (if it's ever deployed)
- fine in theory, but where's the practice? [3] [4]



[3] Briscoe “Designing for tussle; case studies in control over control ” (2004)  
<http://www.cs.ucl.ac.uk/staff/B.Briscoe/present.html#0406pgnet>


[4] Communications Futures Programme & Communications Research Network  
<<http://cfp.mit.edu/>> <<http://www.communicationsresearch.net/>>



# heterogeneity = multiple architectures?

## heroic tussle or pathetic indecision?

- yes, at architecture design time
- yes for testbeds
- but, a spin-off from testbeds for real-life run-time? Please, no!



- for connected internetwork flows and routes must traverse all architectures
- inter-architecture resource control? routing?
- can't even solve these problems for one inter-domain architecture
- do we hear end-customers & app developers saying "If only we had multiple architectures"?

# implications for testbed design

- overlays not useful for e2e resource control expts
  - fine if focusing purely on naming, addressing, routing
  - care! architecture research will eventually need to be integrated
- traditional view of infrastructure testbed problem
  - need real applications, real users
- the fault in the Internet is the fault in our expts
  - our assumptions about operators, businesses, info svcs depts
  - we need real operators, real businesses, real info svcs depts
    - set policies with their own reputations and resources at stake
- the prize is true convergence, 3GPP/IMS, mesh, ISPs, NGNs
  - varying outcomes at the same time: ‘design for tussle’

spare slide  
my research agenda

Q&A



# rebalancing research agenda priorities

- global scale asynchronous event messaging
  - short co-ordination /control messages (discovery, notification, synch, config)
  - control/co-ordination for lower layers (config, routing, failures) as well as apps
  - connecting the physical world to the information world – the Internet of things
  - overlay multicast not panacea for state scaling & many other problems [1]
- resource allocation / congestion control / fairness
  - longest lasting architectural vacuum – becoming acute
  - flow equality goal (TCP) root cause of many problems [2]
    - solutions [3] have been obscured by this dogma
  - hi acceleration for hi-speed short flows



[1] Briscoe “The Implications of Pervasive Computing on Network Design” (2006)

[2] Briscoe “Flow rate fairness: Dismantling a religion” (Oct 2006)

[3] Briscoe et al “Re-feedback and re-ECN”

<<http://www.cs.ucl.ac.uk/staff/B.Briscoe/pubs.html>>

## in summary

- eat your vegetables then you can have your dessert
  - have as much spice as you want on your vegetables
  - classic distributed computing problems to solve
- avoid sexy research fashions
  - active networks, multihop wireless, p2p overlays
  - **unless** treated as exemplars of the classic problems
- instead sex up the classic problems with some tussle

[www.cs.ucl.ac.uk/staff/B.Briscoe/](http://www.cs.ucl.ac.uk/staff/B.Briscoe/)

