Networked Systems Research is Irrelevant

Saleem Bhatti
School of Computer Science, University of St Andrews
http://www.cs.st-andrews.ac.uk/~saleem/

Apologies to Rob Pike

- Rob Pike, Bell Labs, Feb 2000
  "Systems Software Research is Irrelevant"
  http://herpolhode.com/rob/utah2000.pdf

Acknowledgements

- To some grumpy, old, men:
  Ran Atkinson, Ken Carlberg, Jon Crowcroft, Steve Hailes, Tristan Henderson, Peter Kirstein, Graham Knight, Søren-Aksel Sørensen
- To a couple of after-dinner meetings:
  Chip Elliot, Dongman Lee, Craig Partridge
- ... and to various symposia (of the classic, Greek kind :-)

Questions, not answers!

- This conference is about the "Future Internet".
- My talk is a (light-hearted) look at some potentially serious **questions** for the research efforts for the "Future Internet".
- I do not know the right answers ...
- I probably have the wrong **questions** ;-)
- The never ending search for "clue" ...

Why do we do research? [1]

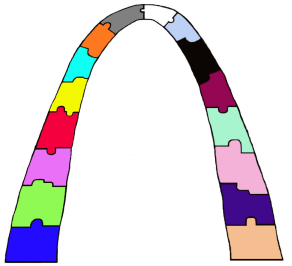

- What will the Future Internet be like?
- None of us knows.
- Each of us has our own ideas of what would be good to have in the future.
- There is a lot of **technology** out there which can (dis)appear quickly.
- We would like to think we have an **architecture** in which to use it all.

Why do we do research? [2]

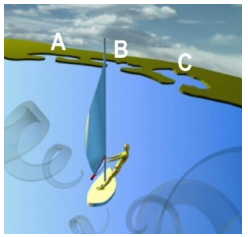

- We are here presenting lots of interesting ideas and work in progress.
- For the Future Internet we would all like our ideas to be used - to be **Deployed**.
- Typically, we would like people to use our ideas because they bring something new to the network - some improvement.
- Lots of other things still to fix:
  - e.g. RFC3869, RFC4948, RFC4984

Some non-functional requirements

1. ***Deployability*****:** Is it easy for us to get our ideas into the network and/or end-systems?
2. ***Backward Compatibility*****:** Can we get our ideas into the network without breaking what is already in the network and/or end systems?

- How important **are** these two issues today?
- How important **should** they be for the future?

Deployability and Backward Compatibility?

- How much do we care about these?
- It seems they are required in some form?
  - Or are they?
  - virtualisation + programmability
- How do we prioritise Deployability and Backward Compatibility?
- How much do we they matter for the Future Internet?

Mobile IP [1]

1) MH arrives at FN, and locates FA (using agent advertisements from FA or by solicitation).

2) MH completes registration procedure with FA.

3) MH updates HA with its new CoA (i.e. the FA).

4) Host A now tries to contact MH. Packets for MH are intercepted by HA.

5) HA tunnels the packets from Host A to the CoA for MH (i.e. the FA)

6) The FA de-encapsulates the inner IP packet and transmits the packet locally to MH.

7) The packets from MH to Host A are sent directly from the FN.

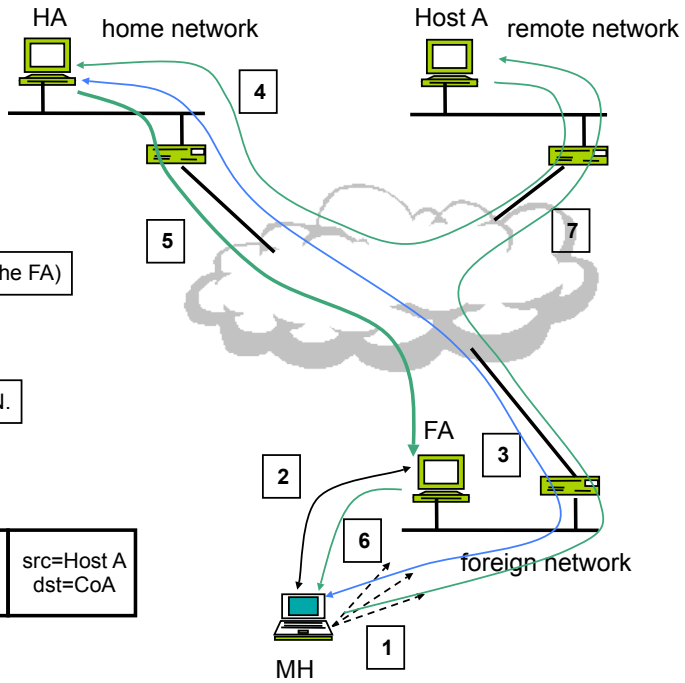

IP-in-IP encapsulation

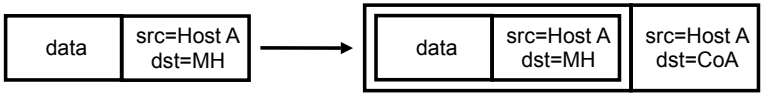

Mobile IP [2]

- ✔ **Transparent to non-mobile hosts**
- ✔ Does not break/change existing IP addressing and routing
- ✔ Can be introduced into the network as required (incrementally)
- ✔ Normal (unicast) routers do not need to be modified
- ✔ Does not affect DNS usage

- ✘ Complex architecture:
  - use of two addresses
  - use of agents
- ✘ Asymmetric routing:
  - could be inefficient
  - TE/QoS
  - higher layer protocol operation (e.g. TCP)
- ✘ Security:
  - firewall configuration
  - authentication
  - end-to-end security
- ✘ No soft hand-off

Mobile IP [3]

- Stateless address auto-configuration:
  - find an address (CoA) for use at the FN
- Neighbour discovery:
  - find default router
- No FA required to support mobility:
  - MH takes care of home address and foreign address
- Need dynamic DNS update support
- **Route optimisation:**
  - **IPv6 Binding Update**
  - **send CoA to remote end-system**
  - **correspondent node *knows* about mobility**
- Security (?):
  - authentication and privacy

Mobile IP [4]

- Mobile IPv4 (MIPv4):
  - not widely implemented or deployed at present
  - complex protocol: mobile node (MN), Home Agent (HA), Foreign Agent (FA)
  - numerous optional optimisations have been proposed
- Mobile IPv4 (MIPv6):
  - also not widely implemented or deployed at present
  - protocol similar to MIPv4 (some optimisations)
  - even more complex with numerous extensions proposed
- IETF MEXT WG

NAT/NAPT [1]

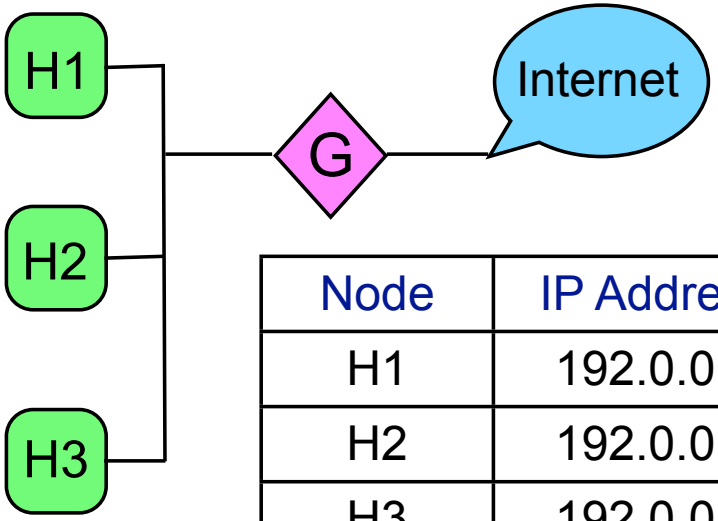

Node	IP Address	Port range
H1	192.0.0.2	5100-5199
H2	192.0.0.3	5200-5299
H3	192.0.0.4	5300-5399
G1	192.0.0.1	5400-5499
G1 (public)	3.1.2.3	-

- G1 uses its 1 public IP address to handle traffic to/ from the Internet for itself and hosts H1, H2, & H3.
- So, G1 is using NAPT and has different TCP/UDP port numbers in public versus on the private LAN segment.

NAT/NAPT [2]

- Some applications use the IP address
  - e.g. FTP uses full IP address
- Rendezvous problems - many sites deploy either NAT or NAPT for perceived security advantages:
  - primarily: remote nodes are blocked from initiating sessions with hosts behind the NAT/NAPT gateway.
  - this can affect some applications (e.g., VoIP).
  - so need a fix for this (e.g., STUN or ALGs)
- Some sites might deploy NAT or NAPT to get IP address portability or to conserve IPv4 addresses:
  - so then engineering fixes are required to deal with the applications that break

Engineering or Architecture?

My definitions for this talk:

- **Engineering**: creating fixes for problems within the current architecture by looking at a focused (narrow) problem space.
- **Architecture**: considering more than just the focused problem space and looking at the design issues of how things work today around that problem space.

How can I test my new systems?

- We are not doing so badly here :-)
  - People generally agree it is good to build and break stuff!
- We have research programmes and testbeds on which to try things out:
  - PlanetLab, Emulab, GENI/FIND, FIRE
- We might even be tempted to share our data and help others to reproduce our experiments ... like the other grown up sciences ;-)

How do we evaluate our new systems?

- What are our evaluation criteria?
- Do we have any metrics that are commonly used?
  - e.g. IPPM WG, TMRG WG
- Where are our shared data sets?
  - e.g. CAIDA, CRAWDAD
- Are these useful for the Future Internet?

Is the end-system stack untouchable?

- Once I have a new system tested, how do I go about getting **real** deployment?
- How do I get this into an end-system stack?
- How many “new” end-system protocols do users have today?
- **Are we simply hostage to the whims of end-system OS vendors?**
- (Meanwhile, the peer-to-peer /overlay chaps are doing rather well ... :-)

Are the network devices untouchable?

- What if I have a new or enhanced network protocol?
- How do I get this into a **real** network device stack on a **real** network?
  - (GENI :-)
- How does anyone get a new/enhanced network protocol deployed?
- **Are we simply hostage to the whims of network device vendors?**

Do users care about any of this?

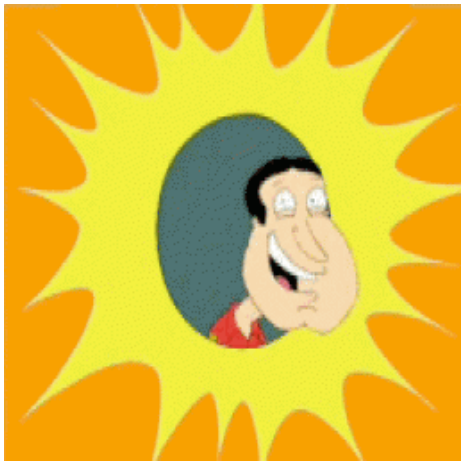

- What about the users?
  - network operators
  - sys/net admins
  - content providers
  - end-users
- How do their needs get considered?
- “Customer pull” could help get our ideas deployed.

Do we have a shared vision?


Bruce Lieberman 50" x 70" oil "Mary" 1983

- For the Future Internet, do we all share the same vision future R&D?
  - Do we need to have a shared vision?
- Do we at least have some understanding and appreciation of the different positions people have taken?
- How do we compare and evaluate our different systems?
- How do we know when we are finished?

Summary

- There are lots of questions (many of them non-technical) we need to ask ourselves for the Future Internet.
- There are likely to be strong non-functional requirements affecting what gets used and deployed.
- Do we know the right questions to ask for research into the Future Internet?