

SWIN  
BUR  
NE

SWINBURNE  
UNIVERSITY OF  
TECHNOLOGY

# Seamlessly integrated broadband services

13 December 2012

Grenville Armitage

[garmitage@swin.edu.au](mailto:garmitage@swin.edu.au)








# I'm surprised the Internet works at all

Many, many millions(+) of loosely co-operating, mostly de-coupled devices....



All trying to locally optimise for their preferred mix of speed,

responsiveness,


consistency,

'quality'

price....




[http://ny.racked.com/uploads/2011\\_11\\_Macys-Black-Friday-Crowd.png](http://ny.racked.com/uploads/2011_11_Macys-Black-Friday-Crowd.png)



# I'm surprised the Internet works at all

Many, many millions(+) of loosely co-operating, mostly de-coupled devices....



All trying to locally optimise for their preferred mix of speed,

responsiveness,


consistency,

'quality'....

price....



What could possibly go wrong?



[http://izismile.com/2009/12/22/how\\_traffic\\_jams\\_start\\_6\\_pics-5.html](http://izismile.com/2009/12/22/how_traffic_jams_start_6_pics-5.html)



# This is not the internet....

---



<http://www.thisiscoossal.com/wp-content/uploads/2012/10/google-3.jpg>



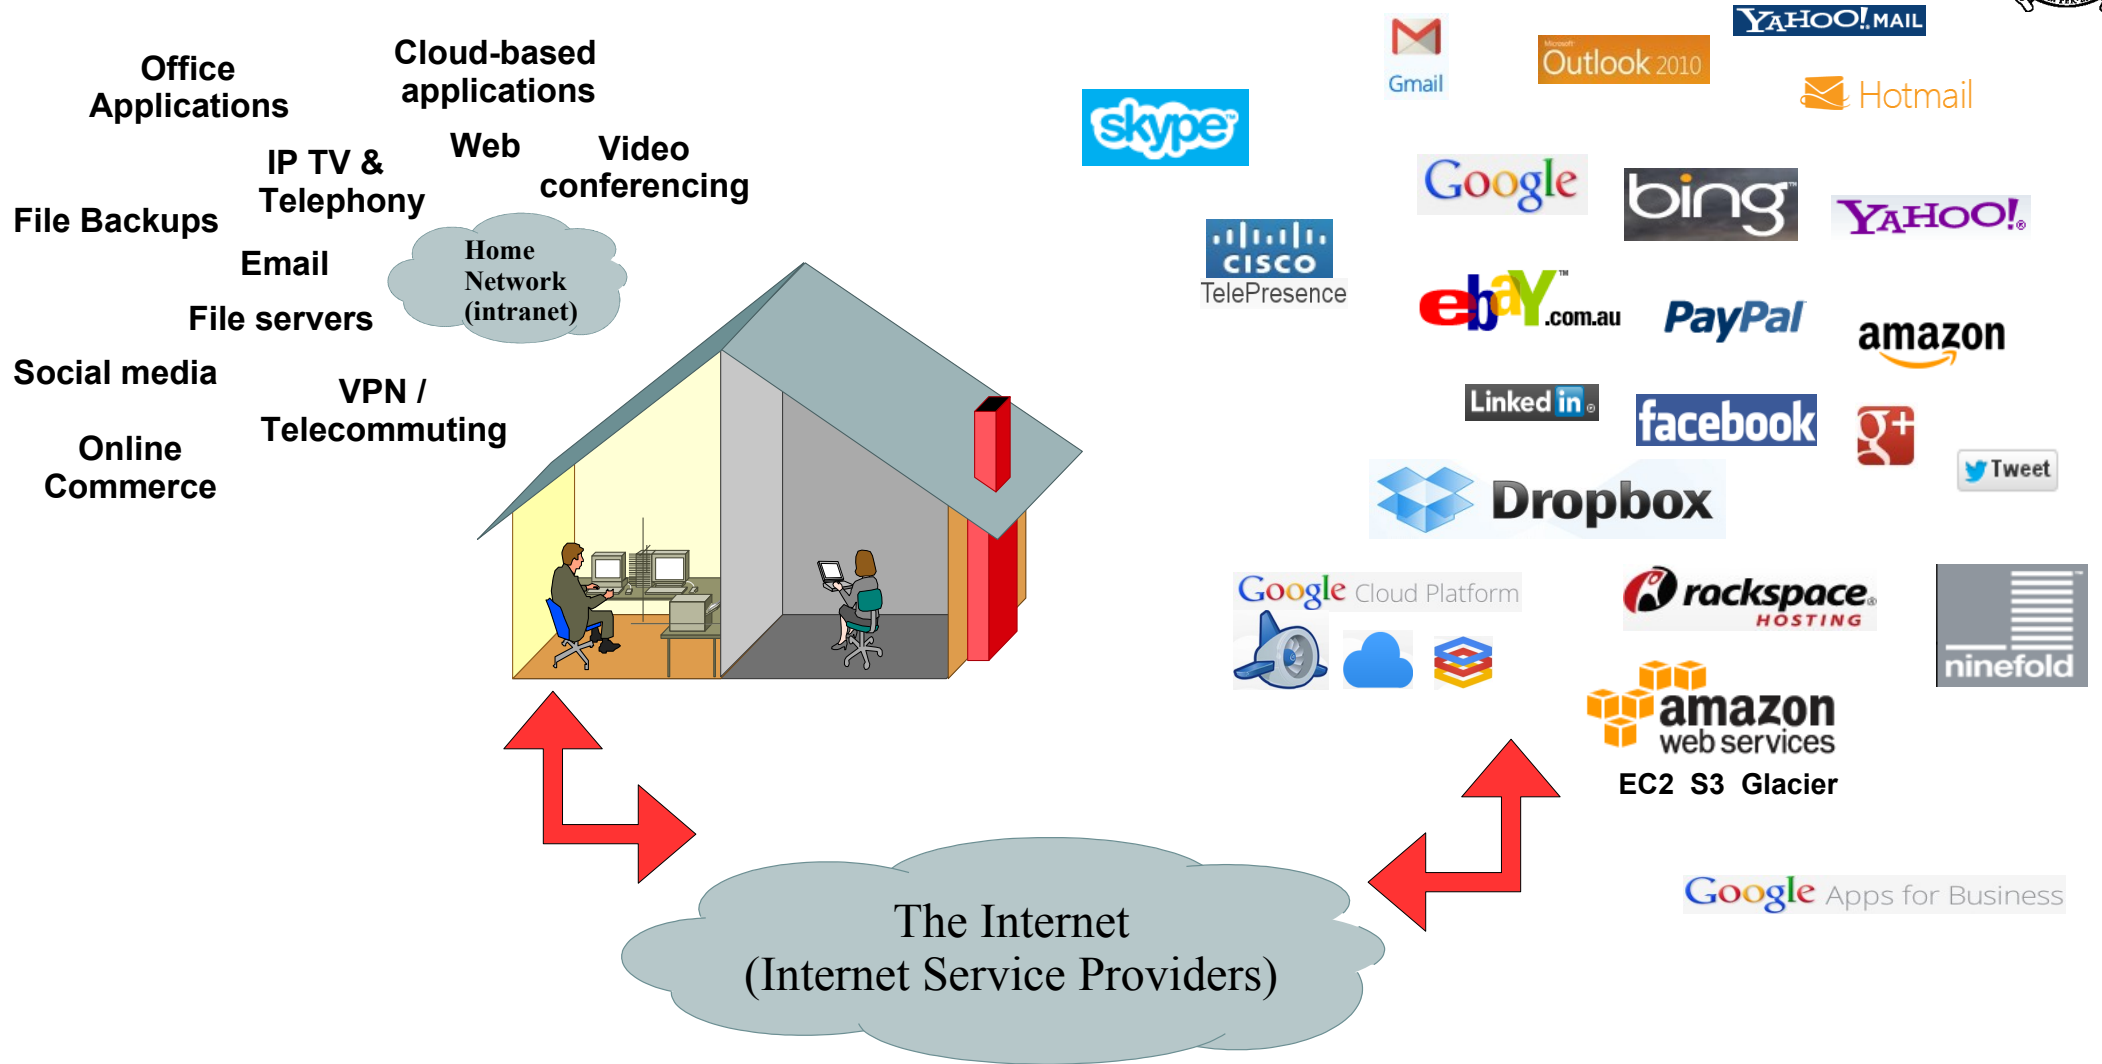
# Getting the feel, but not quite....



[http://hello.eboy.com/eboy/wp-content/uploads/2006/04/FTN\\_CommunicationCity\\_06t.png](http://hello.eboy.com/eboy/wp-content/uploads/2006/04/FTN_CommunicationCity_06t.png)




# Our focus is the plumbing....



(There's a Cloud in here somewhere...)



# When everyone talks to everyone...



The Internet  
(Internet Service Providers)

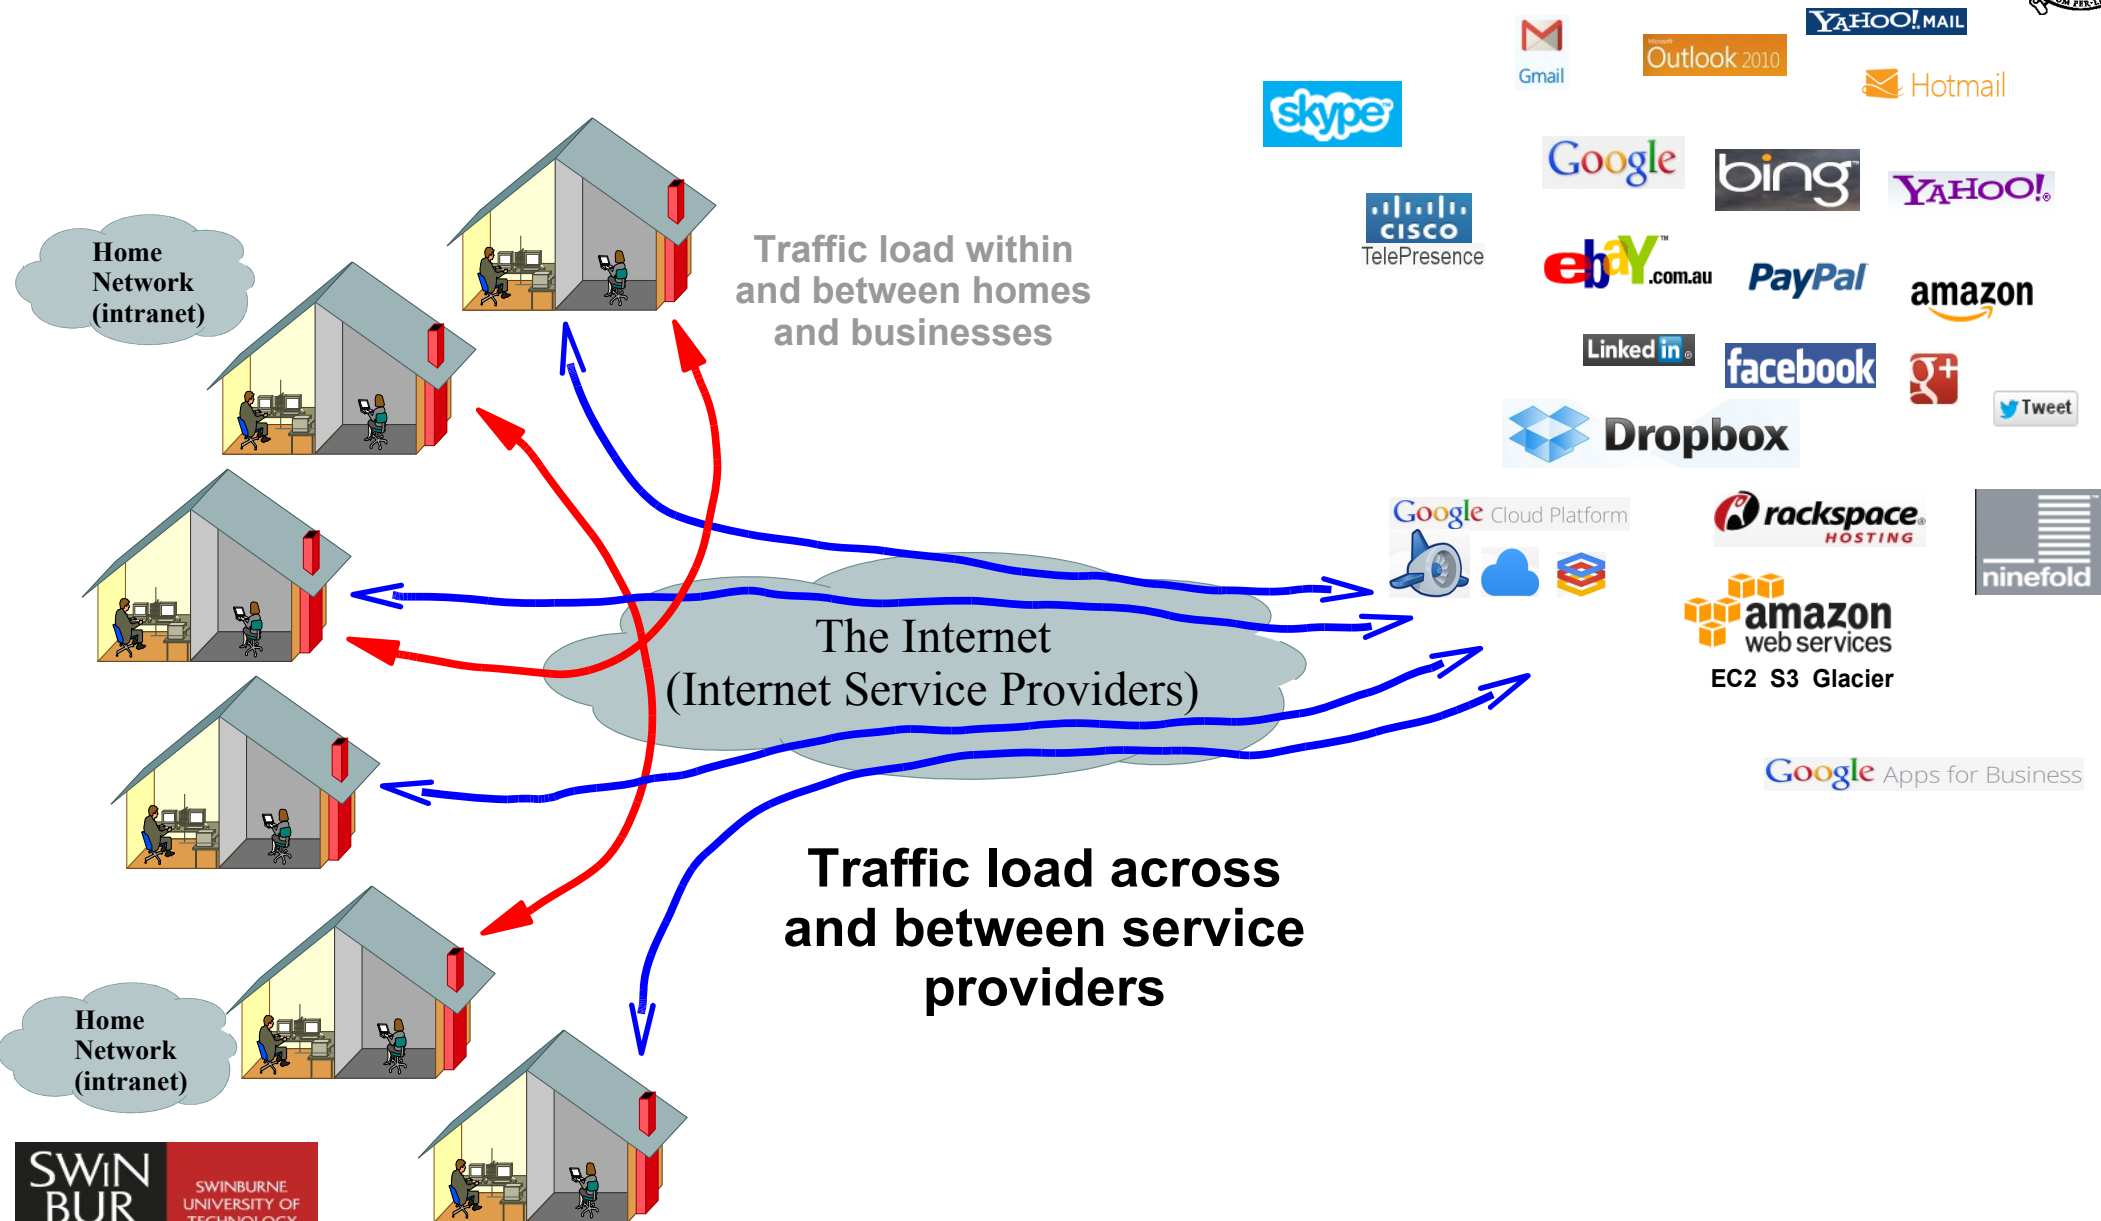
Gmail, Outlook 2010, YAHOO! MAIL, Hotmail, skype, Google, bing, YAHOO!, Cisco TelePresence, eBay.com.au, PayPal, amazon, LinkedIn, facebook, G+, Tweet, Dropbox, Google Cloud Platform, rackspace HOSTING, amazon web services EC2 S3 Glacier, ninefold, Google Apps for Business





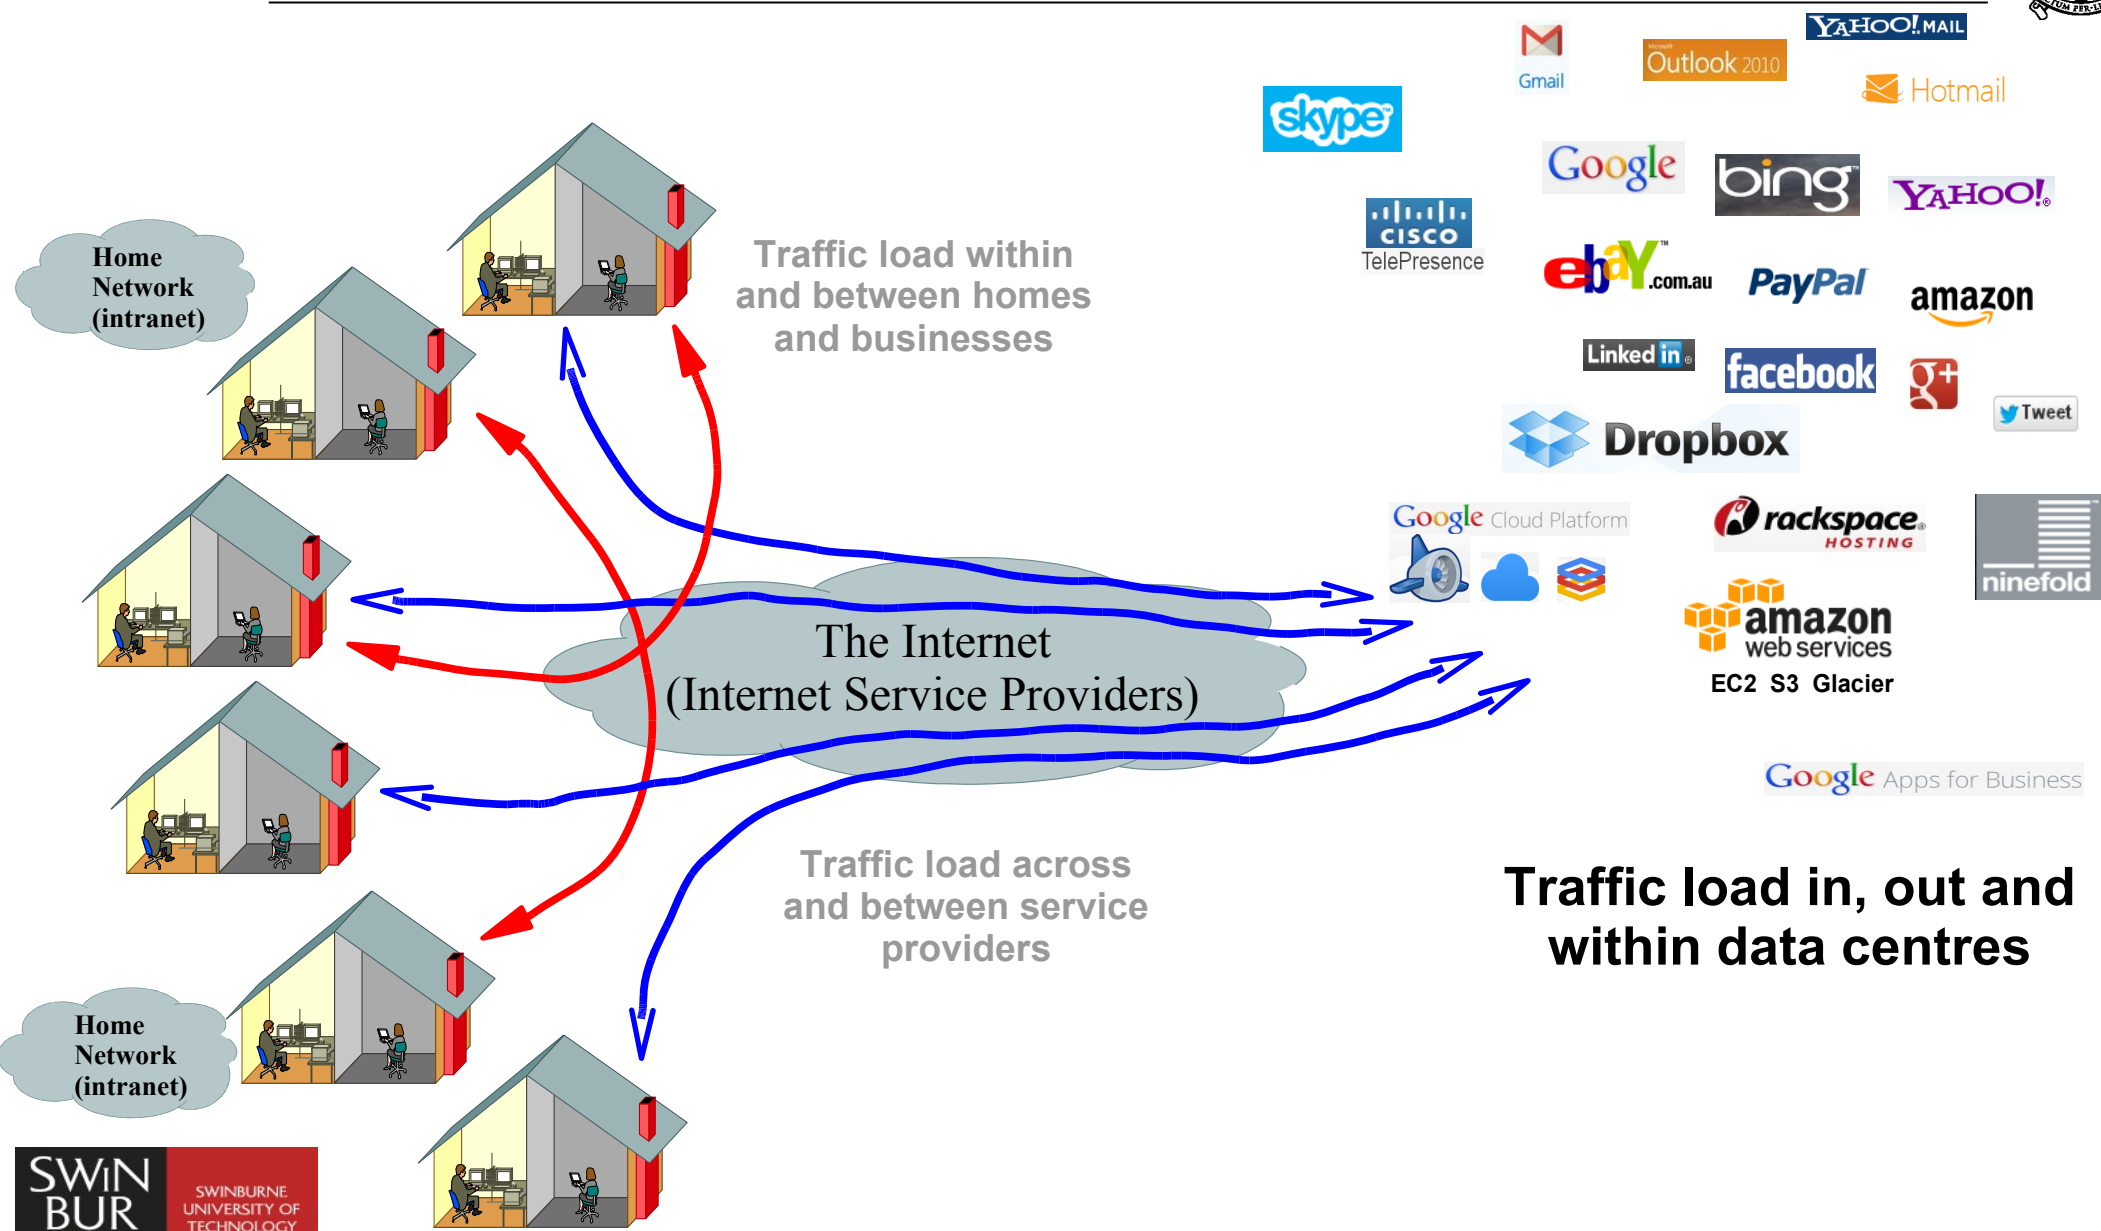


# Packets flow almost everywhere





# Packets flow almost everywhere





# The nutshell of IP networking

---

- Simple packet forwarding service
  - We'll give it our “best effort”
  - Your packet will arrive sometime soon... mostly... or not....
  - Internal queues handle transient overloads
- 'Smart' edge devices
  - TCP – constructs reliable data transfer
  - UDP – provides un-reliable datagrams
- Decentralised / decoupled behaviours
  - Edges infer network conditions indirectly and ~independently
  - Edges recover from forwarding failures (“end to end” principle)
  - Control loops using incomplete, delayed feedback about (knowledge of) network state → fun times!



# “best effort”

---

- Your packet will arrive sometime soon... mostly... or not....

Variable latency

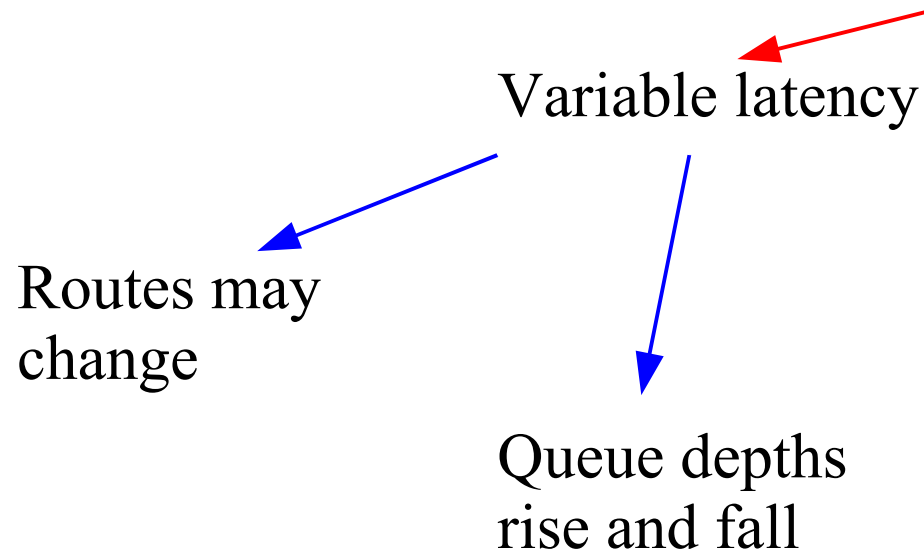





# “best effort”

---

- Your packet will arrive sometime soon... mostly... or not....

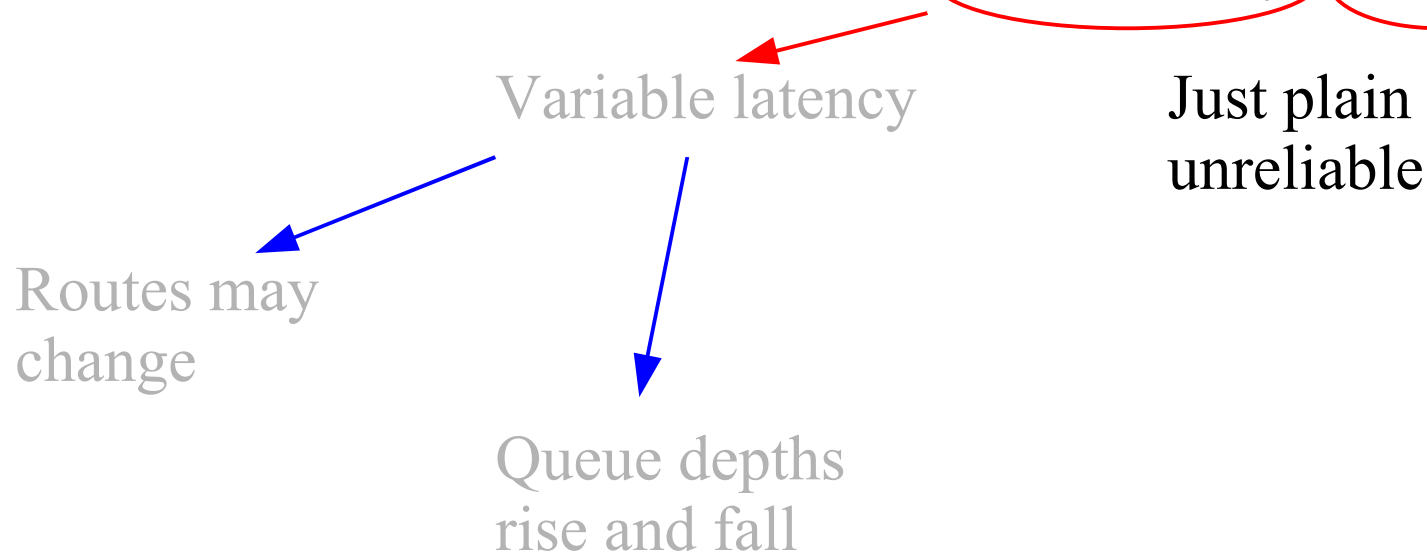




# “best effort”

---

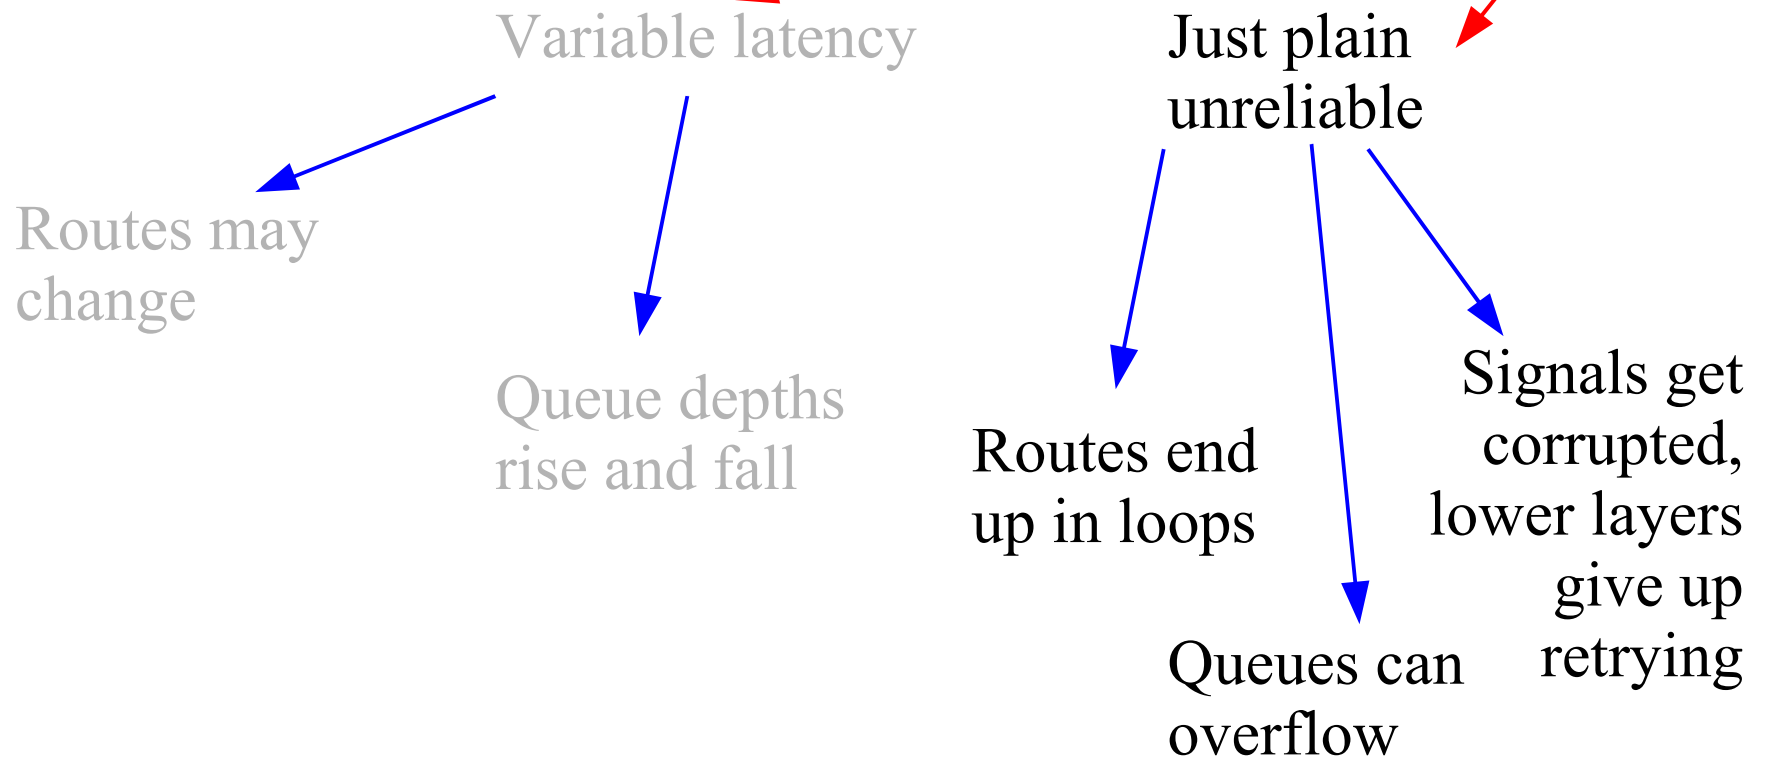
- Your packet will arrive sometime soon... mostly... or not....






# “best effort”

- Your packet will arrive sometime soon... mostly... or not....







---

# The end user and seamless broadband services




# Improving the home user experience

---

- Home gateways are a traffic bottleneck
  - Shared queue → collateral damage
- Understand and model delay-sensitive game traffic?
  - **GENIUS** project (<http://caia.swin.edu.au/genius>)
- Help TCP be a better network citizen?
  - **NewTCP** project (<http://caia.swin.edu.au/urp/newtcp>)
- Automated traffic class isolation?
  - **ANGEL, DIFFUSE** projects (<http://caia.swin.edu.au/urp/diffuse>)
- How much bandwidth do future homes require?
  - **ICE<sup>3</sup>** project (<http://caia.swin.edu.au/ice> & general interest project)

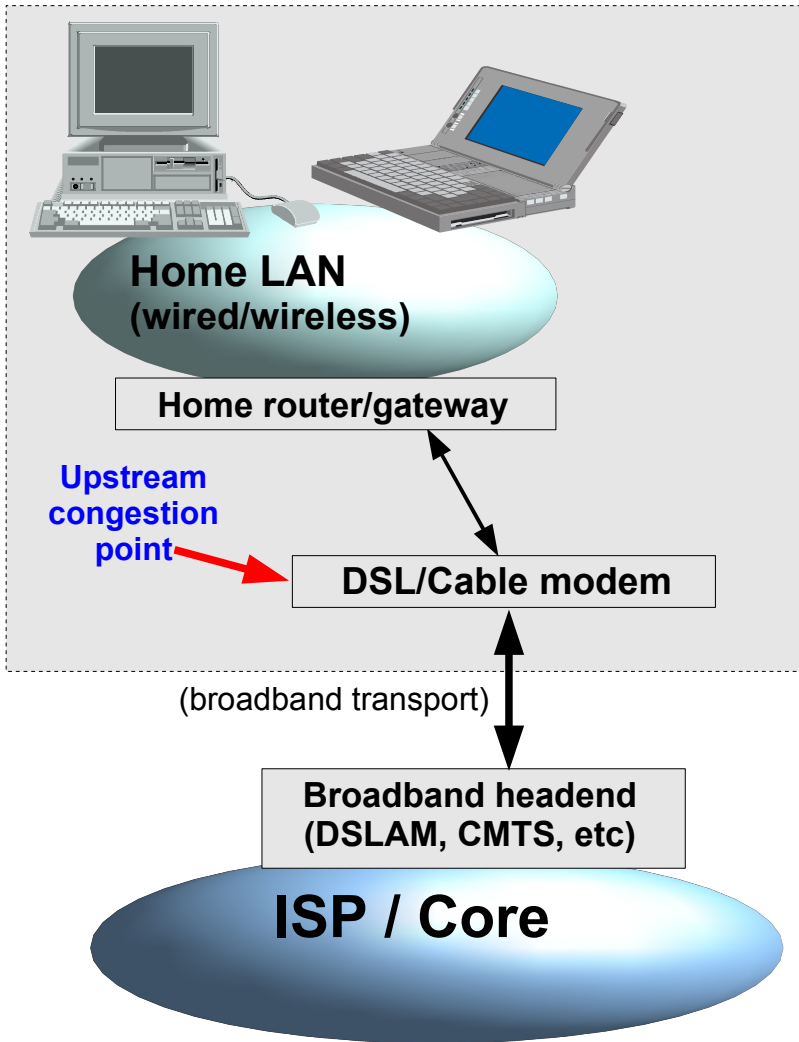


# GENIUS project


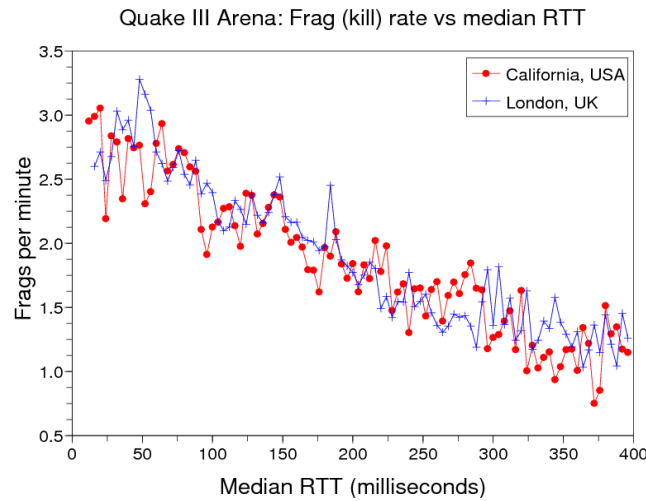




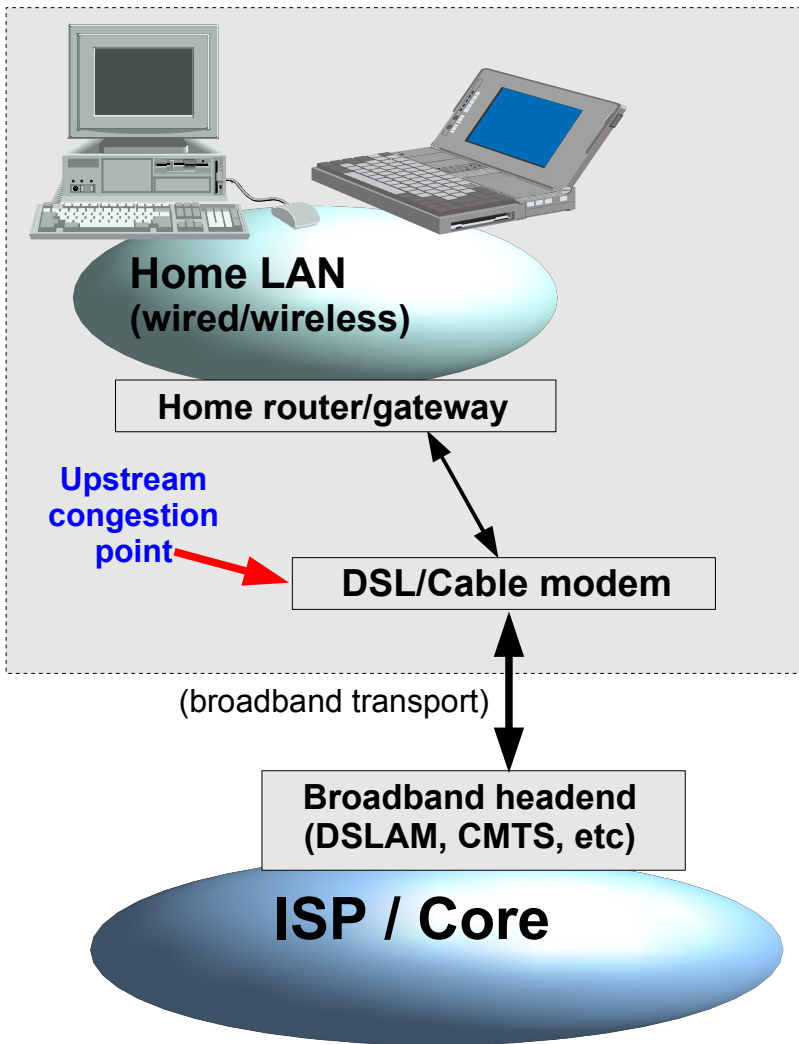
# GENIUS project





## Understand game player tolerance to RTT



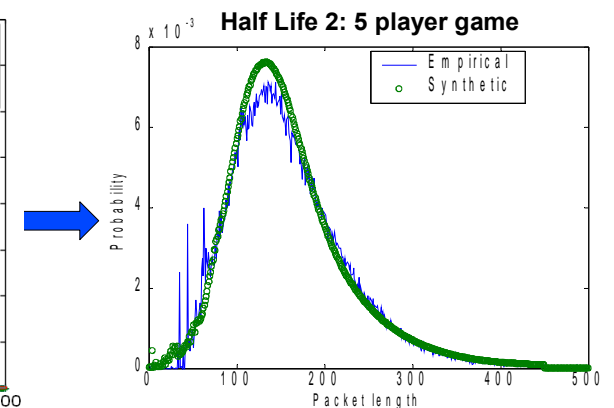

# GENIUS project



## Understand game player tolerance to RTT



Model game traffic patterns through congestion points, produce generator tools....



(Convolution of 2-player and 3-player games)




# NewTCP project

---

- Traditional (*loss-based*) TCP needs packet losses
  - infers network congestion when loss occurs
  - this means someone has filled the bottleneck queue
  - .... which causes additional latency (“collateral damage”)
- CAIA has been exploring *delay-based* TCPs
  - reacts when bottleneck queue *starts* to fill
  - minimises collateral damage to other flows sharing queue
- Also, *multipath* TCP
  - Mobility, and spreading load across multiple interfaces

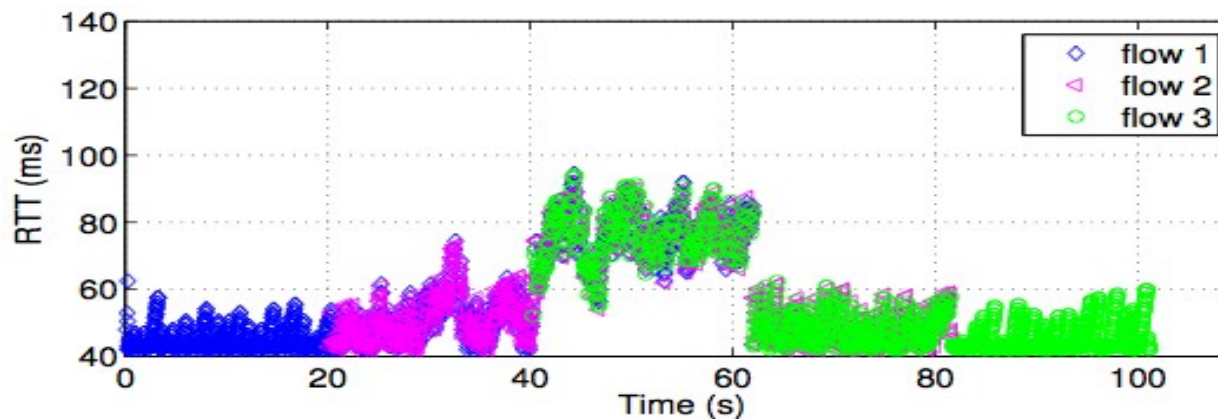


# Example: CAIA's delay-gradient TCP



High/variable induced RTT

(a) NewReno RTT dynamics due to induced queuing delays




Low/modest induced RTT

(b) CDG RTT dynamics due to induced queuing delays


# DIFFUSE project

- A bottleneck with multiple queues *could* isolate TCP and VoIP/game flows
- But identifying VoIP/game flows is beyond most users
- DIFFUSE automates:
  - *Identification* of VoIP/game flows
  - *Prioritisation* of VoIP/game flows
- Minimise TCP's collateral damage → improved user experience

## Diffuse architecture



## Diffuse on a TP-Link home router








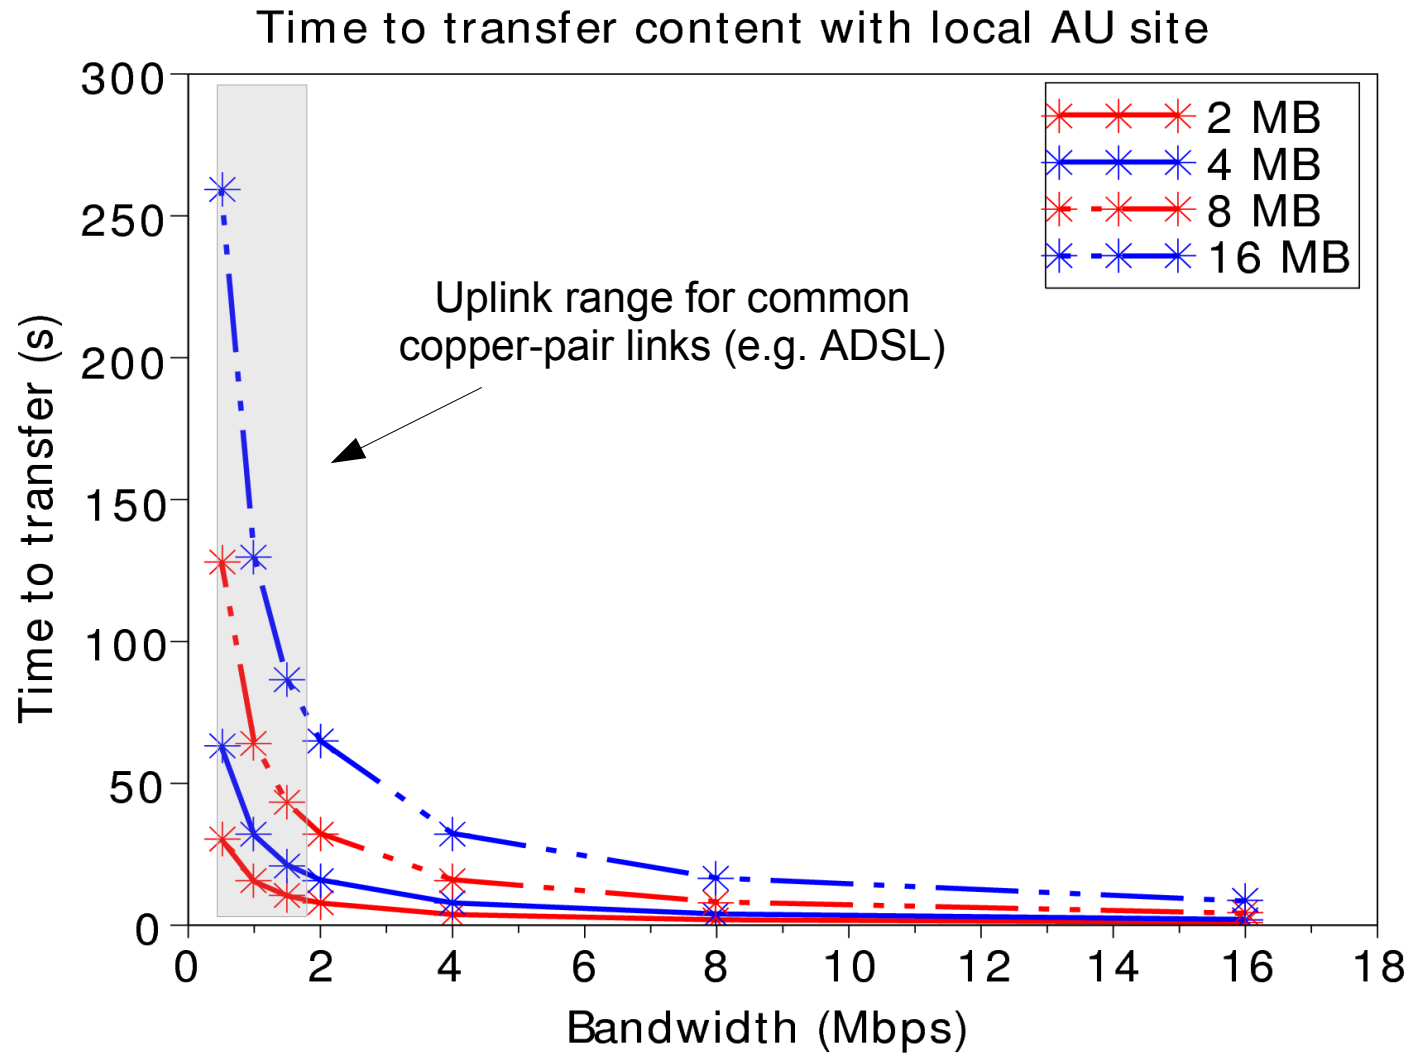
# ICE<sup>3</sup> project

---

- Implications of a densely fibre-connected society
  - More bandwidth at the edge than the core?
  - (NBN?)
- Performance impact of bandwidth vs latency
  - Are there natural or intrinsic limits to bandwidth use?
  - What changes when content is mere milliseconds away?

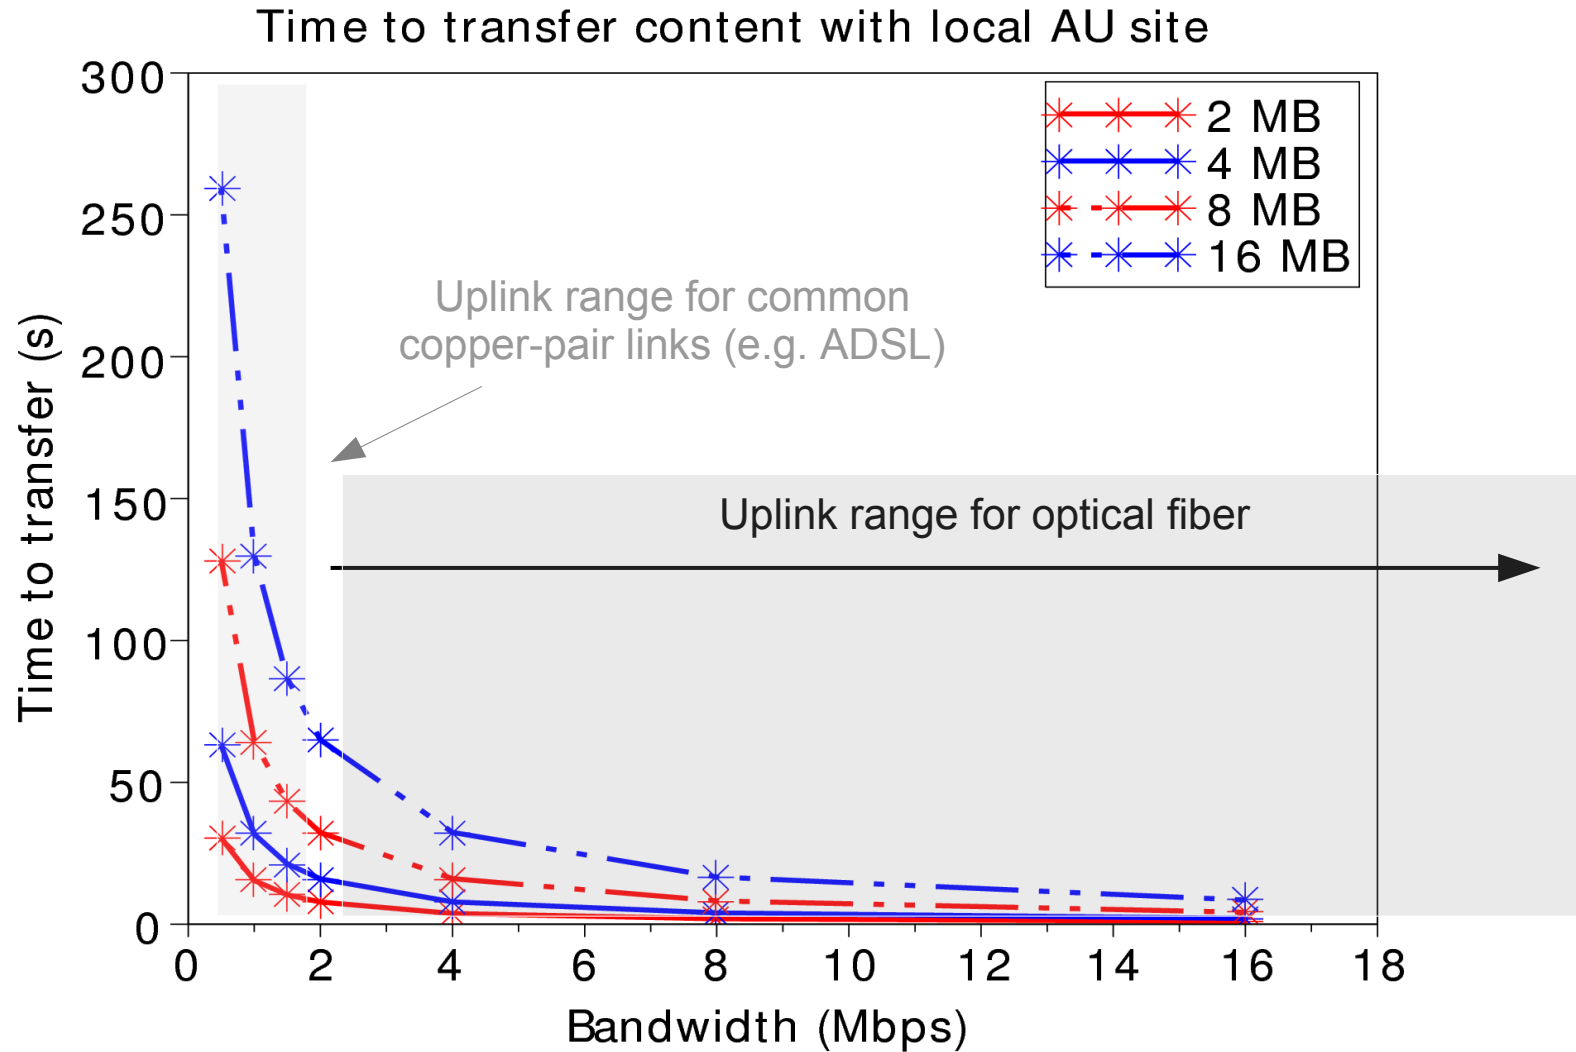


# The need for speed






# The need for speed





# The headache of latency....





---

# Peering inside the network

("You can't manage what you don't measure")



# STING project

---

- Surveying The INternet's Growth (<http://caia.swin.edu.au/sting> )
- Revisiting and enhancing techniques to estimate:
  - ...the fraction of used IPv4 addresses
  - ...current IPv4 address usage patterns
  - ...potential growth in IPv6 address utilisation
- Potential for predicting
  - ...forces acting to promote IPv6 adoption
  - ...market value of remaining IPv4 stock
  - ...future changes to routing tables
  - ...energy consumption envelopes

# STING IPv4 Internet Map

[www.caia.swin.edu.au/sting](http://www.caia.swin.edu.au/sting)

August 2012

Used IPv4 addresses observed from multiple sources between Jan 2011 and Aug 2012

Private, experimental, multicast, reserved

Not allocated

Not routed

80--100%

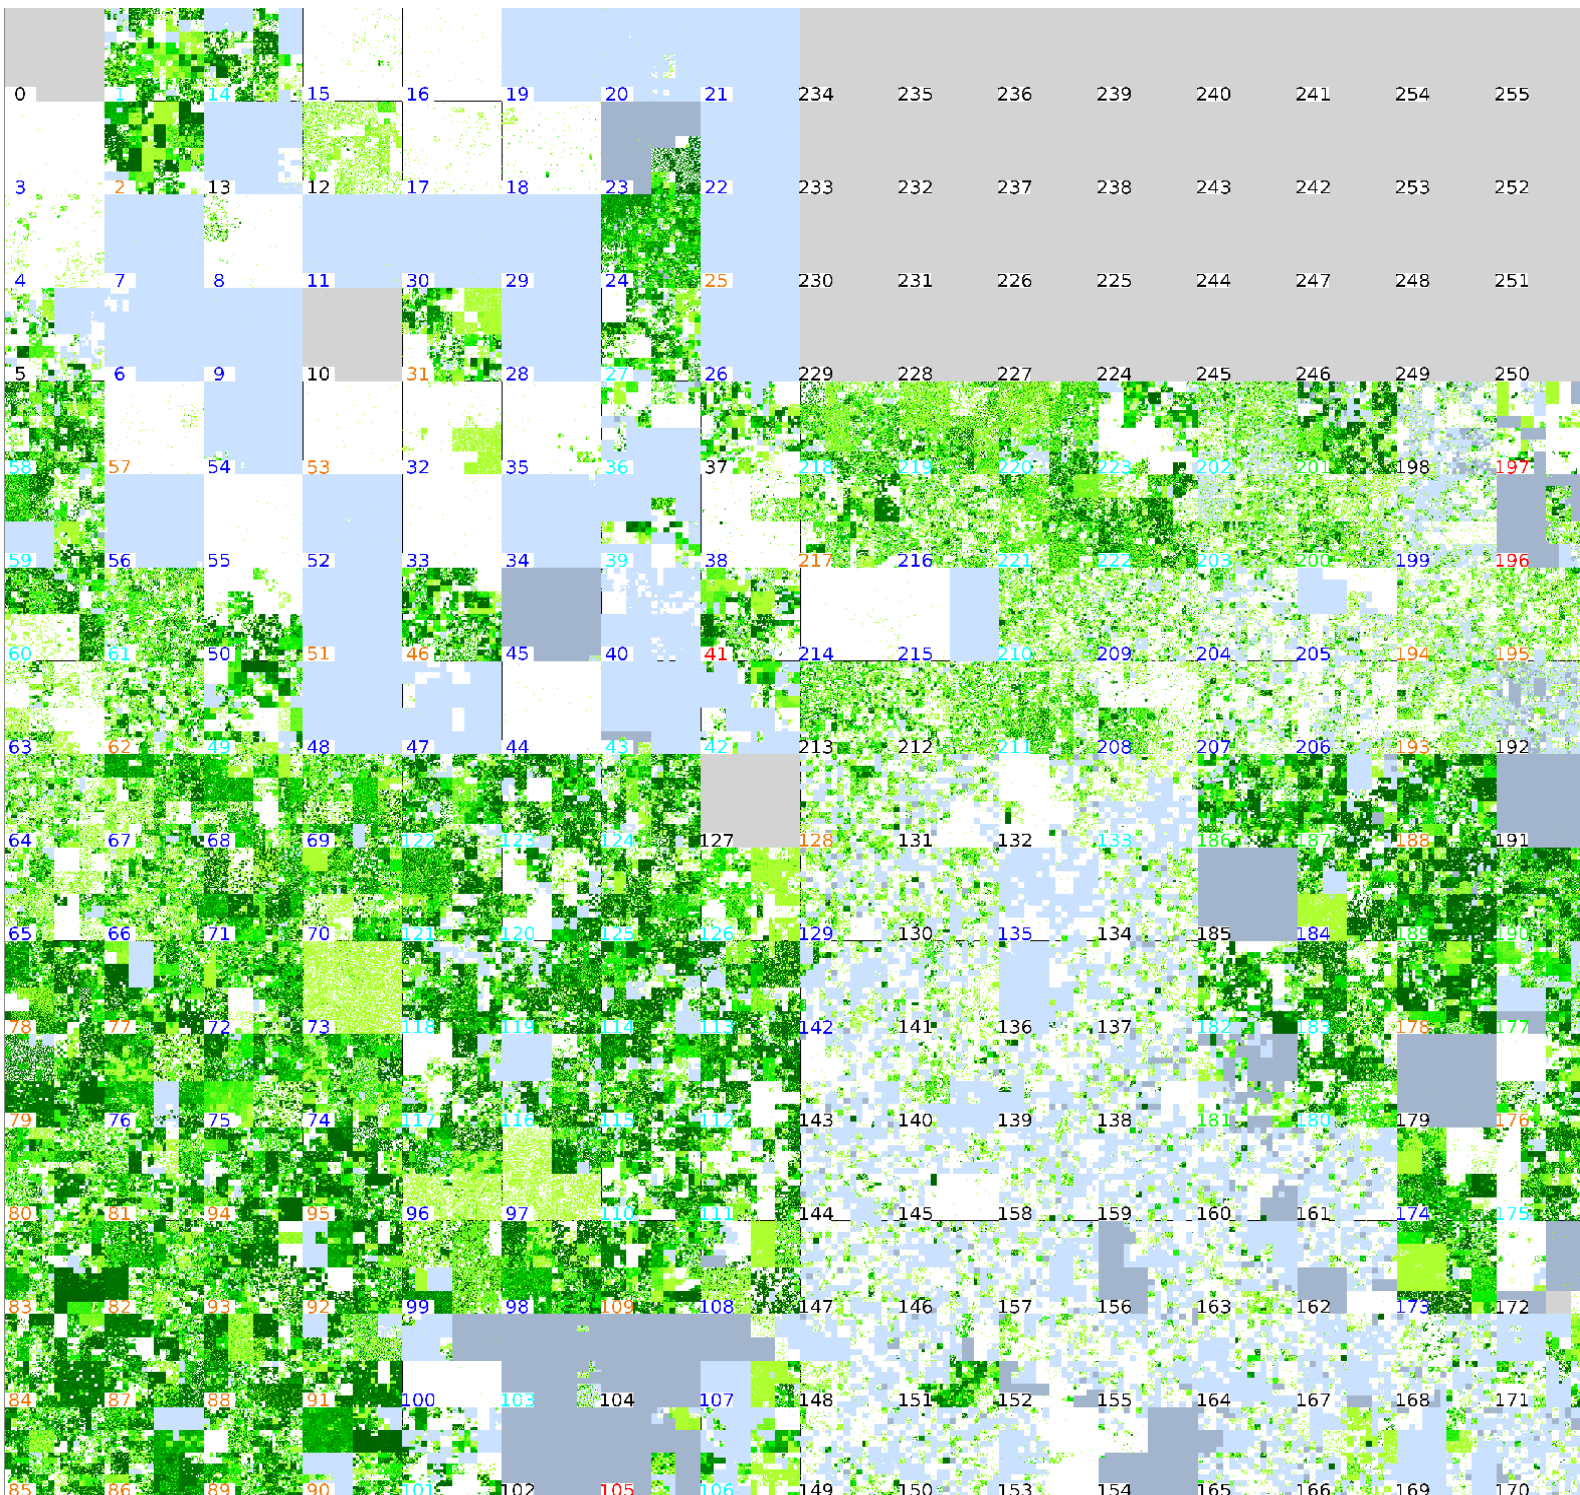

Used IPv4s in /24 subnet

>0--20%

All /8 prefixes are ordered based on a Hilbert curve and numbers are coloured based on the registrar

- AfrinIC
- APNIC
- ARIN
- LACNIC
- RIPE
- Several

Hilbert curve





# Instrumentation

---

- SIFTR (<http://caia.swin.edu.au/urp/newtcp/tools.html>)
  - Detailed 'live' logging of TCP state machines in FreeBSD
  - Enhanced insights into TCP dynamics
  - Used both inside and outside CAIA
  
- SPP (<http://caia.swin.edu.au/tools/spp>)
  - Measures RTT actually experienced by normal traffic
  - No need for synchronised clocks at measurement points
  - Augments many testbed projects @ CAIA






# Routing, Security and monitoring

## ISP / Core



**BGP Heuristics project:** Secure BGP → digital signatures → massive increase in CPU load per BGP update: Find ways to reduce propagation of gratuitous BGP updates!


**Covert channels research:** Can we hide information inside streams of normal IP packets? Why, yes, we can! What are these channel capacities, how can they be defeated....



**Greynets project:** Monitor external probing by monitoring unused IP addresses scattered in amongst used addresses...

**L3DGE project:** Visualise the intrinsically abstract. Real-time monitoring of network state using 3D game engines. Collaborative control of network devices using 'devices' inside the virtual world....





---

The end (for now....)