The Internet at a Cross-Road
Clean-slate vs. Growing Pains

Roch Guérin
(with input and help from many others)

Dept. Elec. & Sys. Eng.
University of Pennsylvania

Penn
UNIVERSITY of PENNSYLVANIA

This Talk

- Like any other technological system the Internet needs to evolve, but this has proven increasingly hard and I will argue not primarily for ***<u>technology</u>*** reasons
  - It is *really big*, but that alone does not explain recent difficulties
  - The IPv6 adoption example (and many others)
- This has led to calls for "clean-slate" initiatives
  - The software defined network (SDN) example
- So are clean-slate approaches and SDN in particular the answer?
  - They can *help* make changes easier, but they wont *eliminate* the challenges
  - The hard part is in understanding not just *how* new technologies will be adopted, but also *why* and in creating incentives for realizing it
- We need to understand the complex interactions that affect adoption decisions
  - Some representative examples
    - Integrated vs. separate networks
    - A "simple" dumb vs. smart network example

Sizing-up the Internet

- The “three” Internet stake-holders
  - Internet end-users (consumers of Internet services)
    - ~117 millions in 1997, ~360 millions in 2000, and ~2.2 billions in 2011 (from ~2% to ~33% of the world’s population)
  - Registered Internet domains (providers of Internet content and services)
    - ~15,000 in 1992, ~27 millions in 2000, and ~138 millions in March 2012
  - Internet Autonomous Systems (providers of Internet connectivity)
    - ~5,000 ASes in 1996, ~10,000 ASes in 2000, and ~60,000 ASes in mid-2012
- Some Internet metrics
  - Core Internet routing tables
    - ~5,000 entries in 1992, ~70,000 entries in 2000, and ~430,000 entries in mid-2012
  - Global IP traffic growth
    - ~5 Tera(10^{12})Bytes/month in 1992, ~84 Peta(10^{15})Bytes/month in 2000, ~28 Exa(10^{18})Bytes/month in 2011, and over 15 Zeta(10^{21})Bytes/month predicted in 2016

It is big, still growing, with interacting stake-holders

Sources:

- http://www.internetworldstats.com/stats.htm
- http://www.dailychanges.com/
- http://en.wikipedia.org/wiki/Internet_traffic
- http://en.wikipedia.org/wiki/Global_Internet_usage
- http://en.wikipedia.org/wiki/Landline
- http://www.zooknic.com/Domains/counts.html
- http://bgp.potaroo.net

The IPv6 Migration

If we want to keep growing the Internet, we need more addresses
But, the IANA pool was depleted in February 2011 (no more addresses left!)

RIR	Projected Exhaustion* Date	Remaining /8s in RIR Pool
APNIC:	**19-Apr-2011 (!)**	1.1896
RIPENCC:	**08-Aug-2012**	2.1605
ARIN:	**24-Jun-2013**	4.6591
LACNIC:	**01-Feb-2014**	3.6378
AFRINIC:	**09-Nov-2014**	4.3139

★ Reaches last /8
Source: http://www.potaroo.net/tools/ipv4/index.html

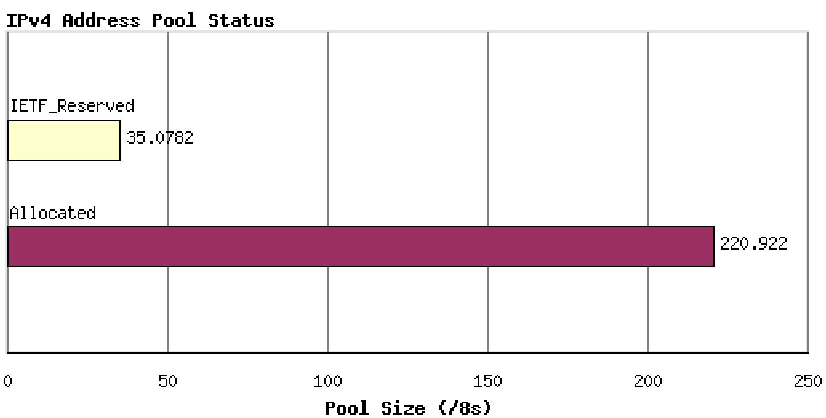

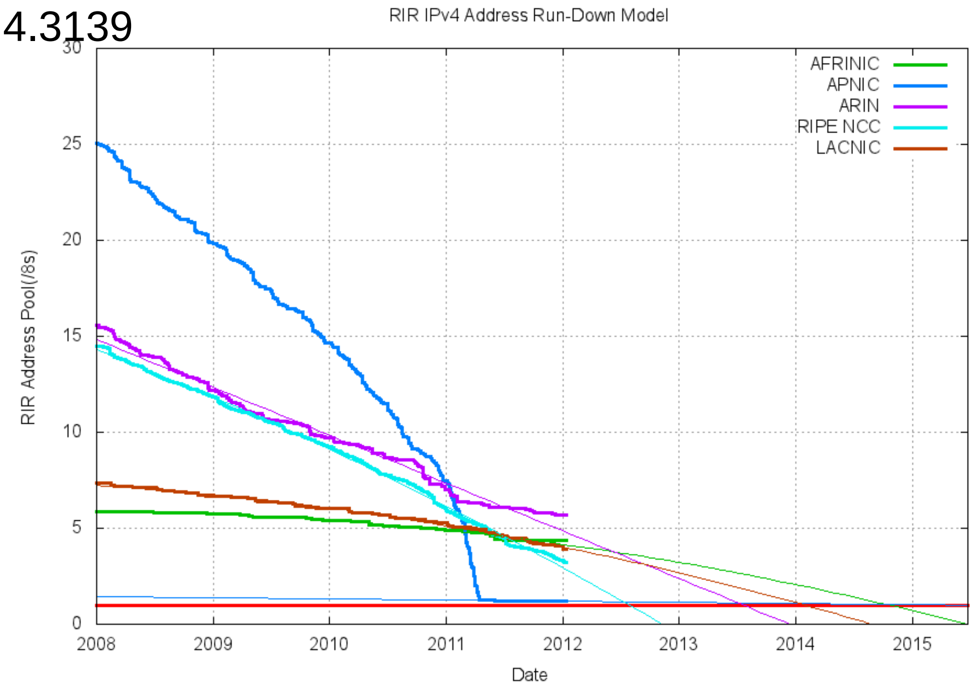

The IPv6 Solution

- We have known about the impending IPv4 address exhaustion for a long time (maybe too long a time)
- The obvious solution, *i.e.*, increasing the address size (IPv6 addresses are 128 bits) was standardized in 1995 (RFC 1883)
- Getting the technology in place took a little time
  - IPv6 has now been systematically available from equipment vendors (routers & hosts) for over 5 years
- So the transition to IPv6 should be relatively straightforward

- Well, not exactly…

The IPv6 “Crisis”

- Current status
  - The problem is real and here to stay (IPv4 address scarcity is not going away)
  - We have a technologically stable solution (IPv6)
  - The solution is facing significant adoption hurdles, including by the major Internet stake-holders
- And it is **<u>not</u>** a technology crisis!

IPv6 Adoption by Content (Web) Providers

(http://mnlab-ipv6.seas.upenn.edu)

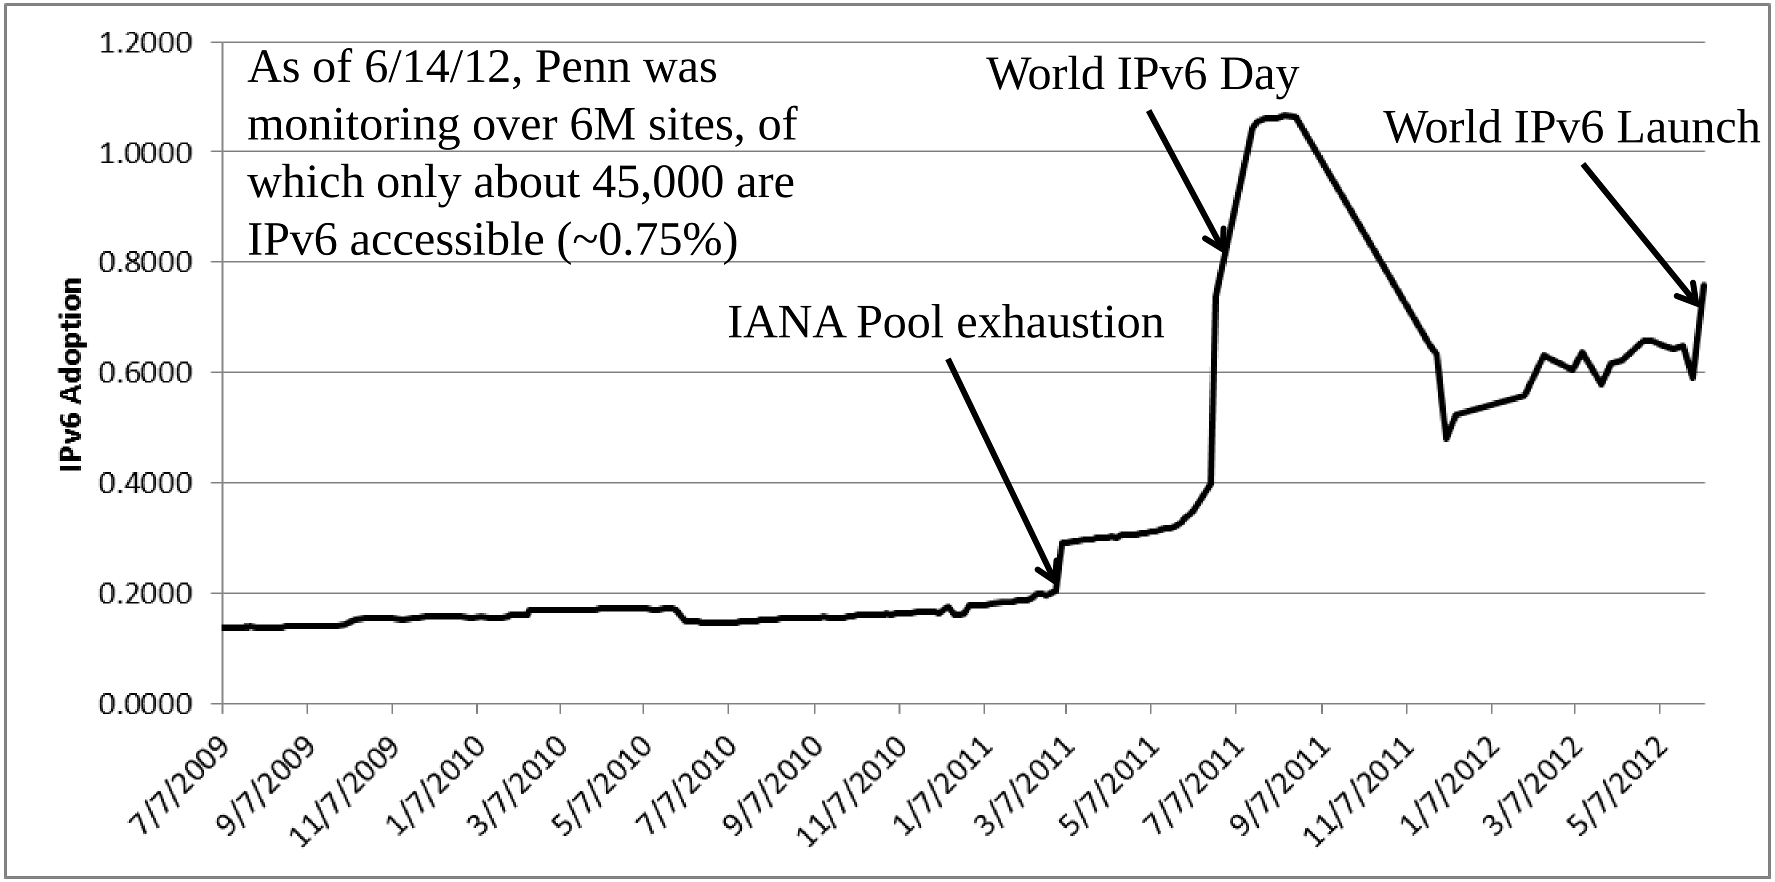

From Content to Service Providers

- Arguably, ISPs should be at the forefront of IPv6 adoption
  - They need more (IPv6) addresses to sign-up new customers and/or address more devices
  - But it is obviously not a sufficient motivation

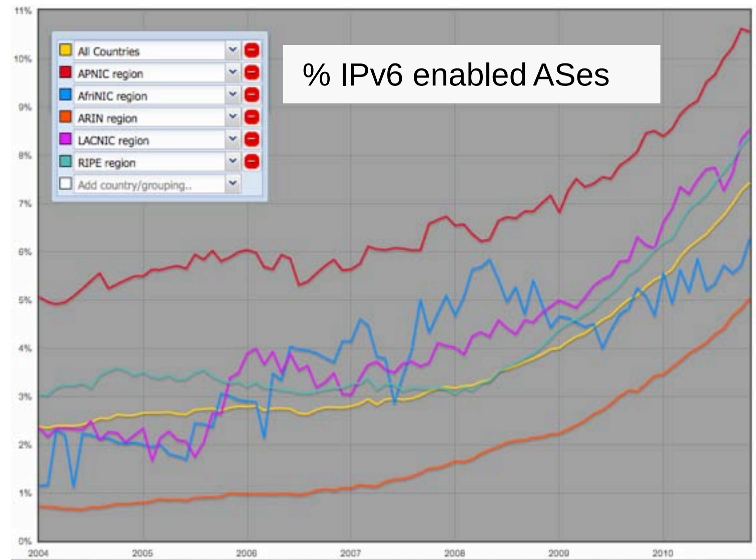

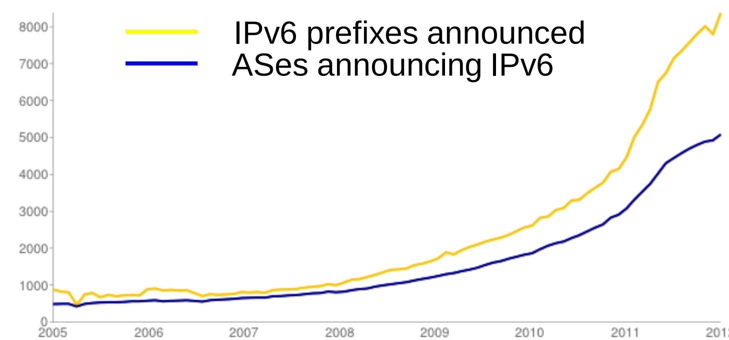

† From http://www.ipv6actnow.org/statistics/

The Problem is <u>**Not**</u> the Technology

- IPv6 systems don't perform worse than their IPv4 counterparts

- A representative example
  - Web access
  - Quantify and explain performance differences

Measurement Vantage Points

Vantage Points		Date on-line	AS_PATH	Type
Comcast	(B)	2/4/11	Y	Commercial
Loughborough U.	(D)	4/29/11	Y	Academic
Penn	(A)	7/22/09	Y	Academic
UPC Broadband	(C)	2/28/11	Y	Commercial
Go6-Slovenia	(E)	5/19/11	N	Commercial
Tsinghua U.	(F)	3/22/11	N	Academic

Measurement Data Overview

- Multiple vantage-points that each
  - Target top 1M web sites (from Alexa) and a few others
  - Record download speeds for all web sites accessible over both IPv6 and IPv4
  - Gather monitoring data over several months
  - Compare IPv6 and IPv4 AS_PATHs
- Clean-up monitoring statistics
  - Confidence targets for individual monitoring rounds
  - Confidence targets for site performance across monitoring rounds (average out temporal variations)
  - Sites that fail to meet confidence targets are eliminated

Quantifying monitoring coverage

Vantage Points	# unique IPs
Comcast	844,355
Loughborough U.	883,413
Penn	1,633,606
UPC Broadband	946,977
Go6-Slovenia	850,954
Tsinghua U.	917,582

Measurement Data Scope

Meet confidence target

# IPv6+IPv4	Comcast	LU	Penn	UPCB	All
Sites (total)	4,568	5,069	12,385	7,843	-
Sites (kept)	3,525	3,906	7,994	4,418	-
Dest. ASes (IPv4)	724	801	1,047	766	1,364
Dest. ASes (IPv6)	592	642	727	609	1,010
ASes crossed (IPv4)	922	1,019	1,332	988	1,785
ASes crossed (IPv6)	742	764	849	746	1,208

P.S.: Removing sites that did not meet confidence targets did not introduce noticeable bias

IPv6 vs. IPv4 Web Access

(when you can, *i.e.*, only ~0.75% of the time)

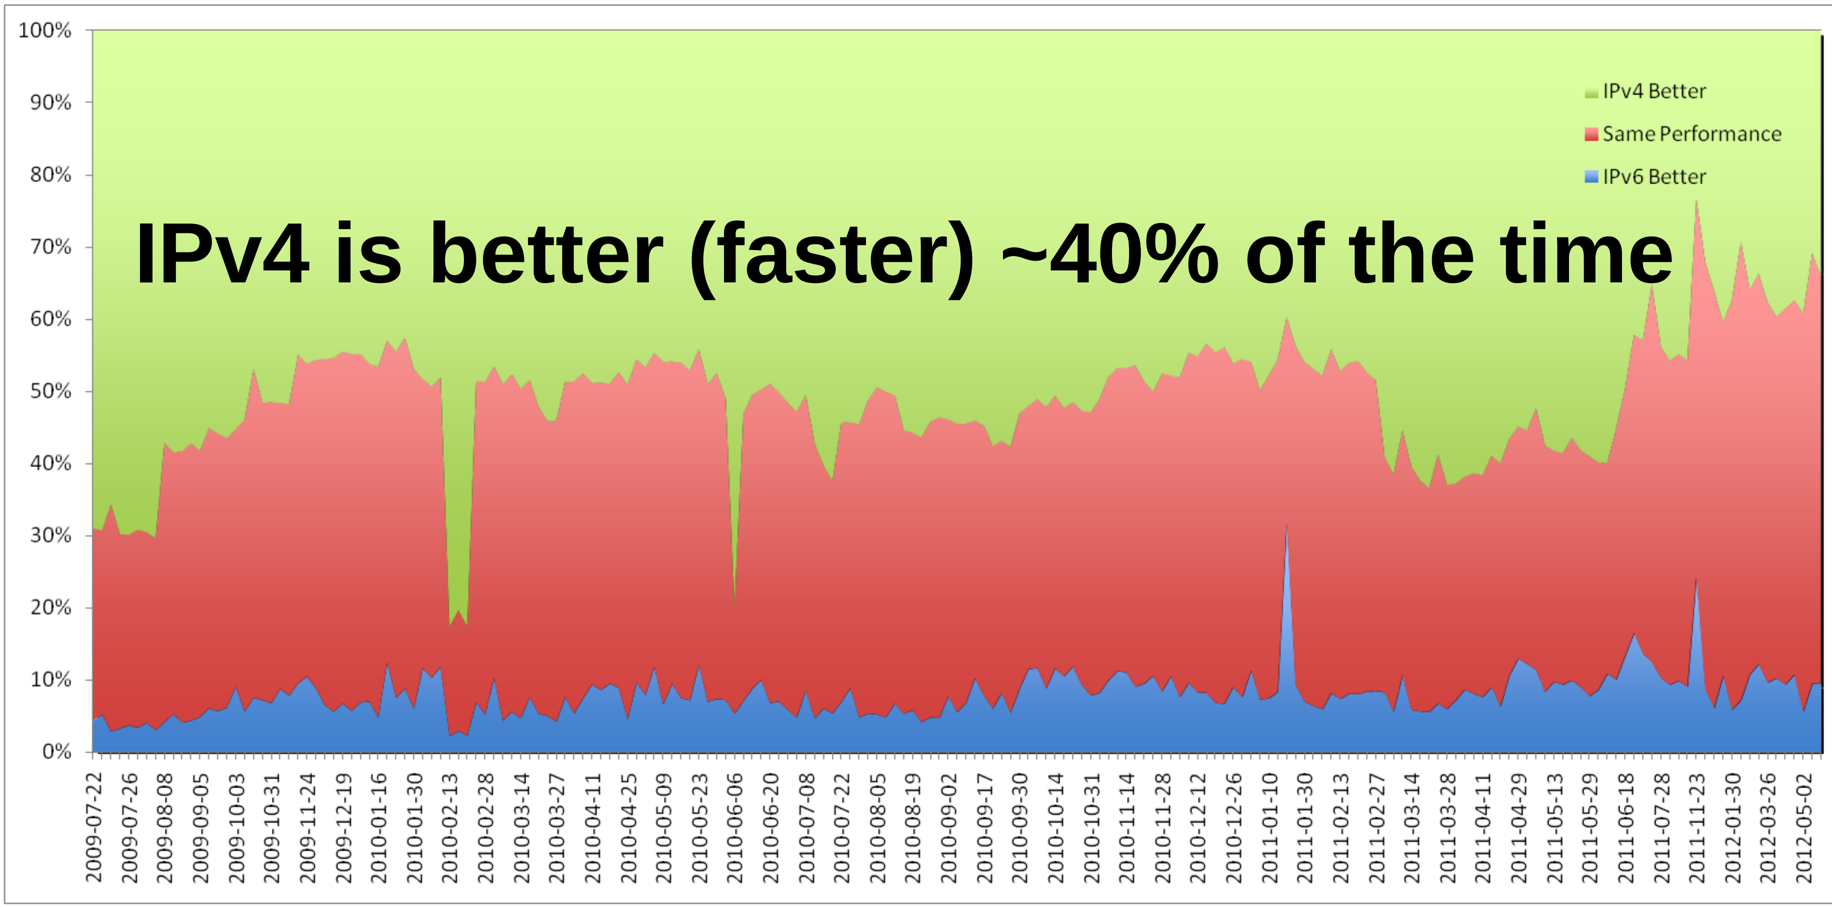

Why Do We See Performance Differences?

- Four major factors can affect how IPv6 performs compared to IPv4

  (E) The client **E**nd-system

  (S) The **S**erver end-system and its access network

  (D) The network **D**ata plane

  (C) The network **C**ontrol plane

Technology

ISP Decisions

Cause and Effect?

Proving Technology's Innocence

- Classify web sites based on whether or not their IPv6 and IPv4 ***"locations"*** and ***"paths"*** differ
  - Same (different) location, *i.e.*, SL (DL) ≡ Same (different) destination AS
  - Same (different) path, *i.e.*, SP (DP) ≡ Same (different) AS_PATH

    (<u>For SL sites</u> only – DL sites have obviously different AS_PATHs)

- For SP sites, **(C)** is absent and so cannot be the culprit
  - But if there are no problems when **(C)** is not there, it makes it a likely suspect
- Are the results different when we consider DP sites?
  - Differences are likely caused by **(C)**

IPv6 vs. IPv4 Performance Within SP

When paths are identical, IPv6 and IPv4 perform similarly

	Comcast	LU	Penn	UPCB
IPv6 ≈ IPv4*	80.7%	70.2%	81.3%	79.8%
Zero mode	6%	10.8%	9.4%	7.3%
Small # sites	13.3%	19%	9.3%	12.9%
# ASes	233	248	75	124
†Cross-check ☑	129	164	47	82
†Cross-check ☒	0	0	0	0

* IPv6 ≈ IPv4: IPv6 performance is within 10% confidence interval of IPv4 performance, or IPv6 outperforms IPv4

† Cross-checking looks for (in)consistent results for ASes found in the same "category," *i.e.*, SP or DP, from different vantage points

World IPv6 Day (6/08/11) Validation (Sites in SP)

World IPv6 Day IPv6 traffic was significantly higher, *i.e.*, data plane performance was tested more extensively

	LU	Penn	UPCB
IPv6 ≈ IPv4	85.7%	92.3%	72.2%
Other	14.3%	7.7%	27.8%
#ASes	42	13	36
Cross-check ☑	17	8	13

IPv6 vs. IPv4 Performance Within DP

A very different result, when IPv6 and IPv4 follow different paths!

	Comcast	LU	Penn	UPCB
IPv6 ≈ IPv4	11%	10%	3%	8%
Zero mode	5%	3%	12%	6%
# ASes	233	248	75	124

- World IPv6 Day Results

	LU	Penn	UPCB
IPv6 ≈ IPv4 (DP)	48.9%	53.5%	51.0%
#ASes	92	114	102
IPv6 ≈ IPv4 (SP)	85.7%	92.3%	72.2%

Recall SP figures →

Learning from the IPv6 "Mistake"

- Technology maturity was **NOT** the culprit!
  - When ISPs do their part, *i.e.*, routing is the same, IPv6 and IPv4 perform similarly
  - Limited IPv6 adoption has, therefore, other causes
- The ossification argument
  - The Internet's success is its biggest enemy
    - Most changes initially afford limited benefits, and when they become truly needed they face an insurmountable "upgrade" task
  - Overcoming this challenge calls for introducing abstractions that decouple functionality from their implementation, and for facilitating the translation of how the functionality is to be performed
    - Akin to how a compiler translates a high-level programming language into device specific instructions

Towards a Cleaner Slate

Software Defined Networks (SDNs)

- A conceptual model of networks that is reminiscent of modern computer systems
  - Hardware, software, and APIs
- As an over-simplification: The data plane (switches) is decoupled from the control plane (routing), and APIs define how the control plane accesses data plane capabilities
  - The data center example: Commodity switches configured by central controller that specifies routing and load-balancing policies

Are SDNs the Answer?

- Clean abstractions and some level of programmability are definitely useful, ***but***

- The Internet is reasonably modular to start with
  - A layered architectures with well-defined interfaces
  - The intra vs. inter-domain separation

  So this cannot be the only reason

- More importantly, changing or upgrading an infrastructure of the size of the Internet is inherently hard
  - It involves complex interactions across stake-holders
  - Even if SDNs can lower the cost of changes, *i.e.*, facilitate the ***how*** of changes, they do not address the ***why*** for those changes

Why is Migrating to IPv6 So Hard?

- Complex interactions between incentives from different stakeholders
  - Internet Content Providers (ICPs)
    - They derive revenue from users, which depends partly on connectivity quality
    - Converting to IPv6 has a cost (direct or indirect)
  - Internet Service Providers (ISPs)
    - Revenue comes from connecting users (and content providers)
    - Costs for operating the network and deploying translation gateways (IPv4-IPv4 or IPv6-IPv4)
  - Users
    - Connect primarily to access content and services
    - They are sensitive to connectivity cost and quality
- An illustrative "chicken-and-egg" problem
  - What happens once IPv4 addresses have been exhausted?
  - ISPs start giving IPv6 address to new users
    - Users cannot access the bulk of the Internet content (most ICPs are only IPv4 accessible)
    - ISPs deploy translation gateways (IPv6-IPv4)
      - A cost that grows with the volume of translation traffic
    - ISPs want to convince ICPs to become IPv6 accessible (eliminates the need for translation)
    - Gateway quality as a possible control knob for ISPs, but
      - If gateway quality is low, ICPs have incentives to adopt IPv6, but users are (initially) unhappy, *i.e.*, fewer users
      - If gateway quality is high, users are happy, but ICPs have no incentives to adopt IPv6

Another Similar Example

- Migrating from 16-bit to 32-bit AS numbers
  - Initiated in 2006 with a similar motivation as IPv6, *i.e.*, impending resource scarcity
  - An easier migration path, *i.e.*, an incremental deployment and transition scheme
- Where do we stand?
  - A little better than with IPv6, but given how much simpler this is, still far from successful

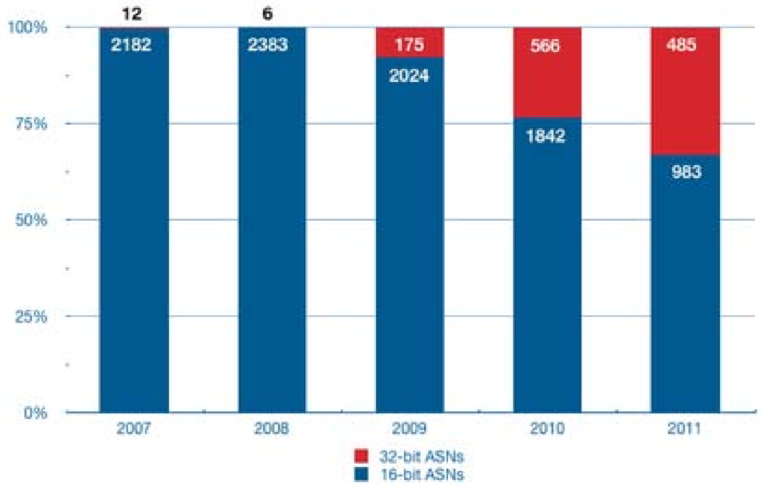

- Since 2007, RIPE has been assigning 32-bit numbers unless a 16-bit number is explicitly requested
  - The vast majority of assigned numbers are still 16-bit
- Why?
  - Over 25% of assigned 32-bit numbers are *returned*, *i.e.*, ASes are unable to get it to work with their provider…

Source: https://labs.ripe.net/Members/mirjam/assigning-32-bit-asns

The Net of It

- Large-scale, complex network systems such as the Internet give rise to a wide-range of interactions that affect technology adoption decisions
  - Both (technology) costs and benefits vary based on the decisions of others

- Understanding those effects and how they impact the deployment of new technologies is as important as the technology itself
  - Tackling those issues calls for an inherently multi-disciplinary approach

Two Illustrative Examples

- Shared or separate networks?
  - When should new services be deployed on an existing (upgraded) network?
- Complex or simple networks?
  - The Internet has popularized the success of the dumb network, but is it always the right answer?

Shared or Separate Networks

- When should we share a network (infrastructure, *e.g.*, cloud) across services vs. deploying them on separate networks?
  - The Internet is arguably a successful example of a *shared* network, *i.e.*, multiple services on the same infrastructure
- Sharing involves both economies and diseconomies of scope
  - Equipment is leveraged across services, but handling multiple services can increase cost and complexity
- Representative examples
  - Triple and quadruple play offerings
  - OT+IT integration in smart buildings

A Simple Attempt

- An existing service (predictable demand) and network
- A new service with uncertain demand
  - Positive demand externalities when services are integrated on the same network
  - Some ability to "reprovision" network capacity in the presence of excess demand (penalty for under-provisioning)
- Economies and diseconomies of scope for both integrated and separate network choices

- Which is better (higher profit):
  1. Upgrade existing network to handle both services
  2. Deploy a separate network dedicated to the new service

Representative Insight

- A wealth of tools from the capacity planning and flexible manufacturing literature
- Identification of two key operational metrics
  - Contribution margin (price less variable costs)
  - Return on capacity (ratio of contribution margin and unit capacity cost)
- More interestingly and less expected, the ability to dynamically reprovision the network can affect the outcome

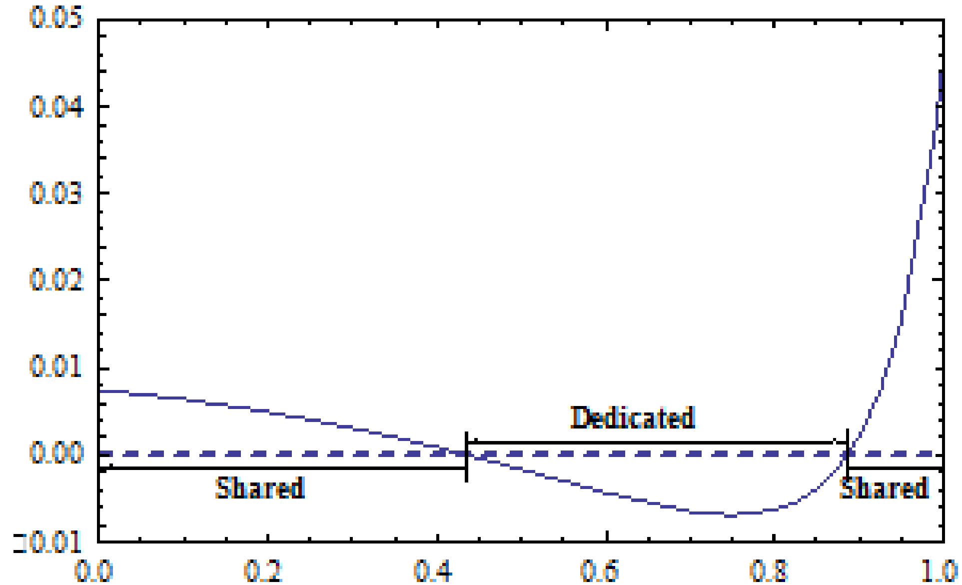

Reprovisioning ability (0-100%)

Difference in expected profits when capacity and demand match exactly

Difference in maximum loss from under-provisioning

The Smart or Dumb Network Question

- The Internet's success has often been attributed to the fact that it is a "dumb" network, *i.e.*, the narrow waist paradigm
- There is, however, a trade-off between the cost and usefulness of adding functionality to the network
  - Dumb networks are cheap, but any added functionality users need has to be developed/bought
  - Smart networks are expensive, but their features can lower the cost of developing new services

Which option is better, when?

Developing Some Insight

- A two-sided market model
  - The network as the “platform”
  - Users and content/application developers as the two sides of the market
- Network
  - Derives revenue from users and developers (more users/developers ⇒ higher revenue)
  - Incurs costs that increase with the number of features it offers
- Application developers
  - Derive revenue from users (more users ⇒ higher revenue)
  - Pay for access to network, and incur development costs that decrease with network features
- Users
  - Pay for access to network
  - Derive utility from accessing applications (more applications ⇒ greater utility)
- What are optimal prices and number of network features?

Representative Insight

- There is an “optimal” number of features
  - The marginal cost increase to the network of adding a feature is equal to the marginal decrease in development cost across developers
- There is not one “right” answer
  - It depends on how network and developer costs are affected by features

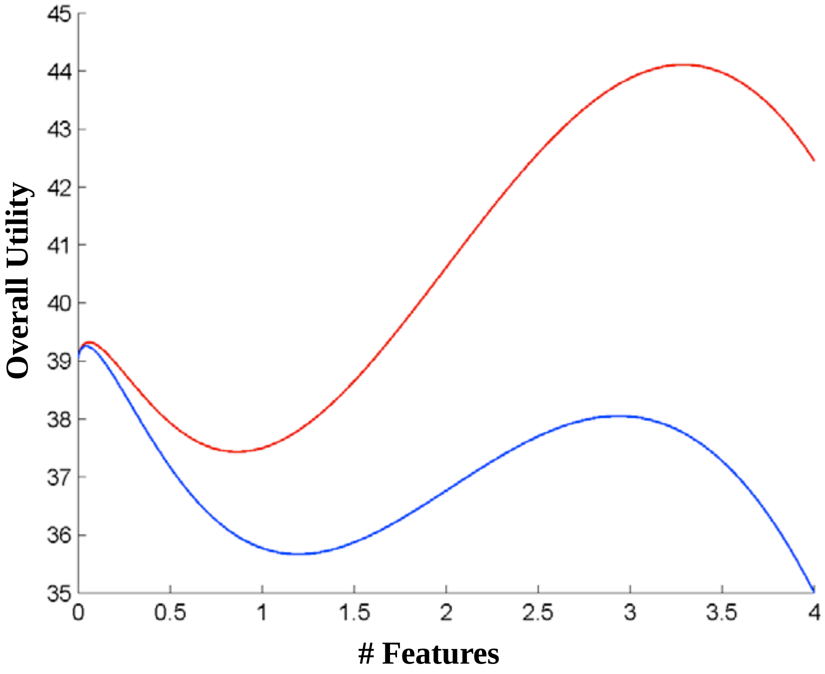

In Summary

- The Internet is
  - One of the largest systems ever engineered
  - A *network* (!) whose "value" is affected in many ways by the decisions of its stake-holders
- This makes predicting the eventual success of new technologies difficult
- Clean (slate) design principles, *e.g.*, SDN, can help by lowering the cost of adding new capabilities and facilitating experimentation
  - But even the best design is evolvable only up to a point
- More importantly, design evolvability is only part of the answer
  - Propagating/understanding changes in large-scale networks is inherently hard
- And there are risks
  - Greater flexibility can result in market fragmentation (Cisco ONE vs OpenFlow)
  - Proliferation of customized approaches can increase operational costs, which can become a significant impediment, *e.g.*, lack of IPv6 operational expertise created deployment hurdles

References & Collaborators

[1]. S. Sen, R. Guerin, and K. Hosanagar, "*Shared Versus Separate Networks – The Impact of Reprovisioning.*" Proc. ACM ReArch'09 Workshop, Rome, Italy, December 2009.

[2]. R. Guerin and K. Hosanagar, "*Fostering IPv6 Migration Through Network Quality Differentials.*" ACM SIGCOMM Computer Communication Review, Vol. 40, No. 3, July 2010.

[3]. S. Sen, Y. Jin, R. Guerin, and K. Hosanagar, "*Modeling the Dynamics of Network Technology Adoption and the Role of Converters,*" IEEE/ACM Transactions on Networking, Vol. 18, No. 6, December 2010.

[4]. M. H. Afrasiabi and R. Guerin, "*Exploring User-Provided Connectivity - A Simple Model*." Proc. ICQT'11 Workshop, Paris, France, October 2011.

[5]. S. Sen, R. Guerin, and K. Hosanagar, "*Functionality-rich versus minimalist platforms: a two-sided market analysis.*" ACM SIGCOMM Computer Communication Review, Vol. 41, No. 5, October 2011.

[6]. M. Nikkhah, R. Guerin, Y. Lee, and R. Woundy, "*Assessing IPv6 Through Web Access – A Measurement Study and Its Findings.*" Proc. ACM CoNEXT 2011 Conference, Tokyo, Japan, December 2011.

[7]. M. H. Afrasiabi and R. Guerin, "*Pricing Strategies for User-Provided Connectivity Services.*" Proc. IEEE INFOCOM 2012 mini-conference, Orlando, FL, March 2012.

Collaborators

- Hadi Afrasiabi, Penn (Ph.D. student)
- Kartik Hosanagar, Wharton, Penn
- Youngmi Jin, KAIST (former post-doc)
- Yiu Lee, Comcast
- Xinxin Li, U. Connecticut
- Mehdi Nikkhah, Penn (Ph.D. student)
- Soumya Sen, post-doc Princeton U. (former Ph.D. student)
- Rich Woundy, Comcast