The Curse of Incompatible (Network) Technologies and the Role of Gateways

Roch Guérin

Dept. Elec. & Sys. Eng
University of Pennsylvania

Boston University, May 7, 2010

Penn
University of Pennsylvania

Acknowledgments

- This is joint work with
  - Youngmi Jin and Soumya Sen (Penn, ESE)
  - Kartik Hosanagar (Penn, Wharton)

- and in collaboration with
  - Andrew Odlyzko (U. Minn)
  - Zhi-Li Zhang (U. Minn)

- References

1. S. Sen, Y. Jin, R. Guerin and K. Hosanagar, "Modeling the Dynamics of Network Technology Adoption and the Role of Converters." Accepted for publications in the IEEE/ACM Trans. Netw. (see http://repository.upenn.edu/ese_papers/496/ for an extended version)
2. R. Guerin and K. Hosanagar, "Fostering IPv6 Migration Through Network Quality Differentials." Submitted for publication. See also http://mnlab-ipv6.seas.upenn.edu/monitor for measurement data.

Outline

- Preamble
  - Common motivations to both parts of this talk
- A model for network technology adoption in the presence of gateways
  - A brief glance at the machinery
  - The insight and surprises
- From "What IPv6 migration?" to "When and How?"
  - A basic model and some associated insights
  - Some preliminary measurement results

Motivations and Approach

- Starting point
  - Over the years there have been many instances of competing network technologies
    - Some have succeeded in displacing the incumbent (SNA $\rightarrow$ IP), some have not (ATM)
  - What affect network adoption decisions?
    - Technology cost and quality, but also externalities
    - Availability of gateways between incompatible technologies
- Building an understanding: A two-prong approach
  1. A generic model for the adoption of incompatible technologies in the presence of gateways
     - Users choose technologies based on a "utility" function
  2. A closer look a IPv6 migration (or lack thereof)
     - A simple model based on quality differentials between IPv4 and IPv6
     - Some initial measurements to assess where we stand

MODELING NETWORK MIGRATION IN THE PRESENCE OF GATEWAYS

Problem Formulation

- Two competing and incompatible technologies
  - Different qualities (q_1 and q_2) and prices (p_1 and p_2)
  - Value of technology also depends on number of adopters (externalities)
    - Tech. 1 is the incumbent (current market penetration of x_1, $0 \leq x_1 \leq 1$)
    - Tech. 2 enters the market with zero initial penetration ($x_2 = 0$)
- Users individually (dis)adopt either technology or none ($0 \leq x_1 + x_2 \leq 1$)
  - Decision based on technology ***utility***
- Gateways/converters enable inter-operability
  - Independently developed by each technology
  - Gateways can be of different quality α_i, $0 \leq \alpha_i \leq 1$, $i = 1,2$
  - Gateways can be of different types and have different properties
    - Duplex vs. simplex (independent in each direction or coupled)
    - Asymmetric vs. symmetric (performance/functionality wise)
    - Constrained vs. unconstrained (performance/functionality wise)

Utility Function

Technology 1: $U_1(\theta, x_1, x_2) = \theta q_1 + (x_1 + \alpha_1 \beta x_2) - p_1$

Technology 2: $U_2(\theta, x_1, x_2) = \theta q_2 + (\beta x_2 + \alpha_2 x_1) - p_2$

- Users are heterogeneous in their valuation of technology (measured by θ)
  - Private information for each user, but known distribution
- Users derive benefits from using technology i
  - Benefits intrinsic to the technology and its quality (θq_i)
    - Tech. 2 better than tech. 1 ($q_2 > q_1$)
  - Externalities grow (linearly) with the number of accessible "users," *e.g.,* Metcalfe's law
    - Normalized to 1 for tech. 1
    - Scaled by β for tech. 2 (possibly different from tech. 1)
    - α_i, $0 \le \alpha_i \le 1$, $i = 1,2$, captures impact of gateways
- Users pay a (recurrent – as with services) price (p_i) for using technology i

Illustrating the Model

1. IPv4 $\leftrightarrow$ IPv6
   – Duplex, asymmetric, ***constrained*** gateways

2. Low def. video conf. $\leftrightarrow$ High def. video conf.
   – Duplex, asymmetric, ***unconstrained*** converters

Most “network” examples involve duplex and often asymmetric converters

IPv4 (Tech. 1) $\leftrightarrow$ IPv6 (Tech. 2)

IPv4: $U_1(\theta, x_1, x_2) = \theta q_1 + (x_1 + \alpha_1 \beta x_2) - p_1$

IPv6: $U_2(\theta, x_1, x_2) = \theta q_2 + (\beta x_2 + \alpha_2 x_1) - p_2$

- Migration drivers
  - IPv4 address depletion (2011-2012)
    - Providers may start giving users the option of a cheap IPv6 address or a more expensive IPv4 address ($p_{IPv4} = p_1 > p_2 = p_{IPv6}$)
  - As “technologies” IPv4 and IPv6 essentially similar ($q_1 \approx q_2$ and $\beta = 1$)
- IPv6<->IPv4 gateways are mandatory
  - Most content is ***not*** yet available on IPv6
    - Few incentives for content providers (more on this later)
  - Gateways are (constrained,) duplex, and asymmetric, and deployed by service providers, *e.g.,* DSL-Lite, etc.
- Users select IPv4 or IPv6 based on price, accessible content (x_1 and x_2), and quality of content access (α_1 and α_2)

Low-def. video $\leftrightarrow$ High-def. video

Low-def: $U_1(\theta, x_1, x_2) = \theta q_1 + (x_1 + \alpha_1 \beta x_2) - p_1$

High-def: $U_2(\theta, x_1, x_2) = \theta q_2 + (\beta x_2 + \alpha_2 x_1) - p_2$

- Service offering
  - Two video-conf service options : Low-def & High-def – different formats
    - Low-def has lower price ($p_1 < p_2$), but lower quality ($q_1 < q_2$)
  - Converters offer interoperability, but video is an asymmetric technology
    - Encoding is hard, decoding is easy
      - Converters allow Low-def subscribers to ***display*** high-def signals but not generate them

    $\Rightarrow$ Both Low-def and High-def subscribers enjoy higher externality benefits when talking to High-def subscribers ($\beta > 1$) – the converse is obviously not true
- Converters characteristics
  - Duplex, asymmetric, unconstrained ($\alpha_1 \beta$ not constrained to be ≤ 1)
- Users choose service that gives them the best overall value
  - Low-def has lower price and can enjoy High-def benefits (if there are High-def users…)

User Technology Valuation

- *Indifference points,* $\theta_1^{\,0}(\underline{x})$, $\theta_2^{\,0}(\underline{x})$, $\theta_2^{\,1}(\underline{x})$, identify decision thresholds for each technology
  - $U_1(\theta, \underline{x}) > 0$ if $\theta \geq \theta_1^{\,0}(\underline{x})$ - Tech. 1 becomes attractive
  - $U_2(\theta, \underline{x}) > 0$ if $\theta \geq \theta_2^{\,0}(\underline{x})$ - Tech. 2 becomes attractive
  - $U_2(\theta, \underline{x}) > U_1(\theta, \underline{x})$ if $\theta \geq \theta_2^{\,1}(x)$ - Tech. 2 over Tech. 1

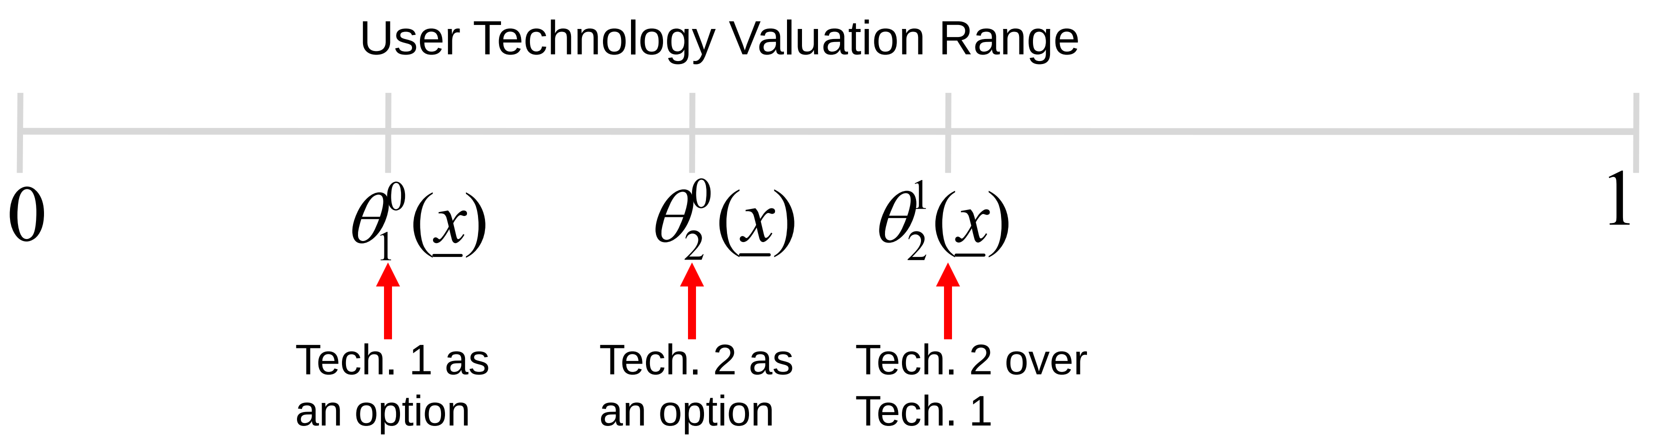

User Decisions

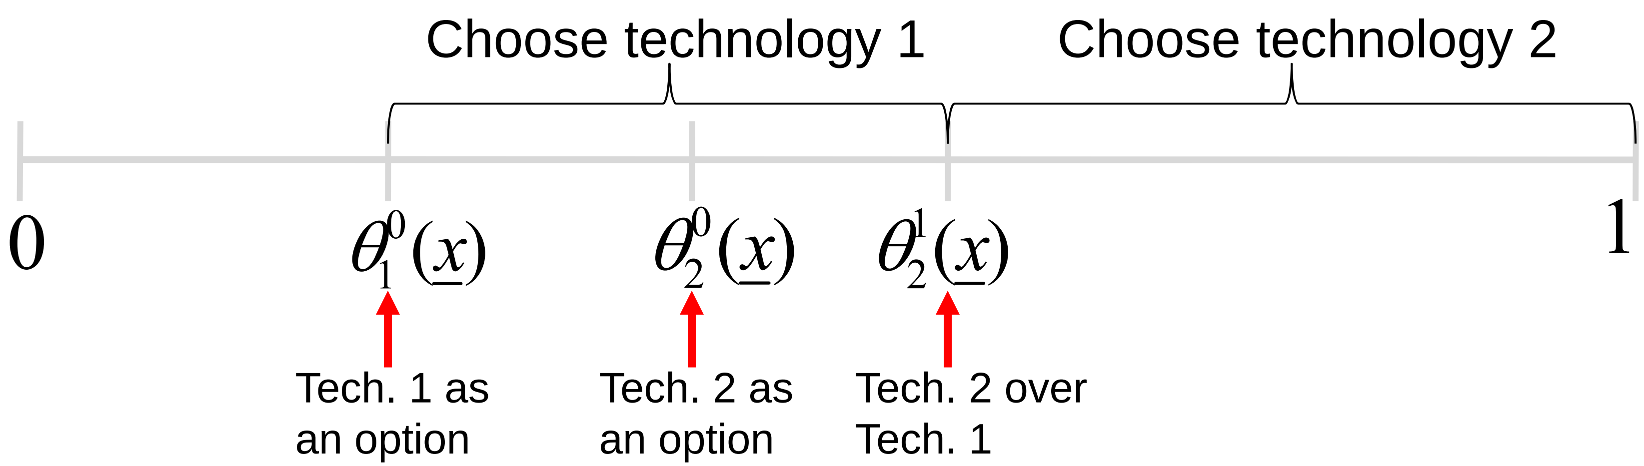

- Users make rational and incentive compatible decisions
  - No Technology if $U_1 < 0,\ U_2 < 0$
  - Technology 1 if $U_1 > 0,\ U_1 > U_2$
  - Technology 2 if $U_2 > 0,\ U_1 < U_2$
- Decisions can/will change as $\underline{x}$ evolves ("switching" costs, *e.g.,* contract breaking penalties, can be added)

Modeling User Decisions

- Let $\underline{x}(t) = \left(x_1(t), x_2(t)\right)$ denote adoption levels at time t
  - This then implies an *hypothetical* number of users, $H_i(\underline{x}(t))$, who should adopt technology i at time t
    - In other words, given adoption levels at time t, we can compute what they should be at time $t+\Delta$
  - At ***equilibrium***, adoption levels should satisfy $H_i(\underline{x}^*) = x_i^*,\ i \in \{1,2\}$

- From hypothetical to actual decisions: **Adoption dynamics**
  - Not all users learn about the current penetration levels at the same time (information diffuses)
  - Not all users react instantly to information about new penetration levels (rate of adoption in target population)
  - Modeling approach: A diffusion process with constant rate $\gamma < 1$

$$\frac{dx_i(t)}{dt} = \gamma\left(H_i\left(\underline{x}(t)\right) - x_i(t)\right),\ i \in \{1,2\}$$

Solving the Model

- It's cumbersome because the indifference thresholds introduce boundaries that demarcate regions with different functional expressions for adoption dynamics, and trajectories can cross region boundaries

- But it is solvable and we can compute/characterize
  - All combinations of possible stable (and unstable) equilibria
  - Global adoption trajectories (concatenation of trajectories across regions)

Adoption Trajectories

- Trajectories in each region are of the form

$$x_i(t) = a + be^{-\lambda_1 t} + ce^{-\lambda_2 t}, i \in \{1,2\}$$

  where λ_1 and λ_2 can be positive, negative, or even complex

- Trajectories can be stitched together as they cross region boundaries
  – Continuous functions with continuous derivatives

Key Findings – (1)

1. The system can have at most two stable equilibria (combinations of only 2 out of 3 among Tech. 1 wins, Tech.2 wins, and Tech. 1 and Tech. 2 coexist)
   – When multiple equilibria are feasible, initial penetration determines the outcome (basins of attraction of each equilibrium)
2. Gateways can help either technology
   – Technology 2 can only benefit from better gateways, while they can harm technology 1
3. Better gateways can harm overall penetration

A "Typical" Outcome

- Separatrix passes through unstable equilibrium and delineates basins of attraction of each stable equilibrium

- Final outcome is hard to predict from the initial evolution of adoption

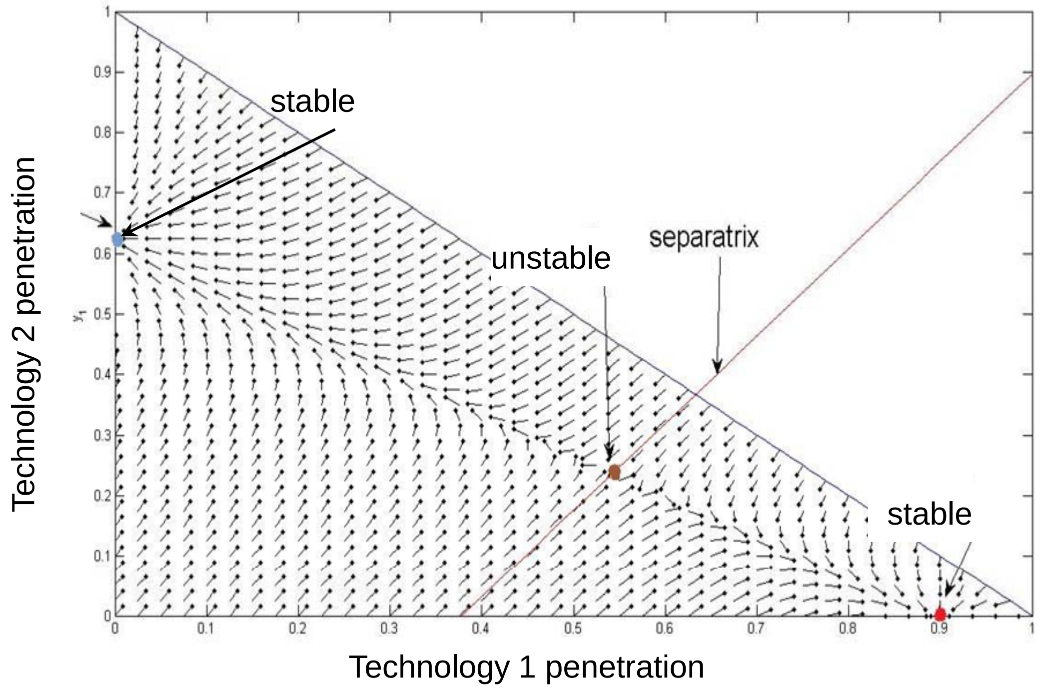

Key Findings – (1)

1. The system can have at most two stable equilibria (combinations of 2 out of 3 among Tech. 1 wins, Tech.2 wins, and Tech. 1 and Tech. 2 coexist)
   – When multiple equilibria are feasible, initial penetration determines the outcome (basins of attraction of each equilibrium)
2. Gateways can help either technology
   – Tech. 2 always benefits from better gateways, while better gateways can occasionally harm Tech. 1
3. Better gateways can harm overall penetration

Gateways Help the IPv6 Entrant

- Assumes IPv6 cheaper than IPv4 and mostly equivalent quality wise ($q_{\text{IPv6}} \approx q_{\text{IPv4}}$)
- In the absence of gateways, IPv6 never takes off unless IPv4 initial penetration is very low, which it isn't…
- After introducing gateways, IPv6 eventually takes over, irrespective of IPv4 initial penetration
  - There is a "threshold" value (**70%**) for gateway efficiency below which this does not happen!

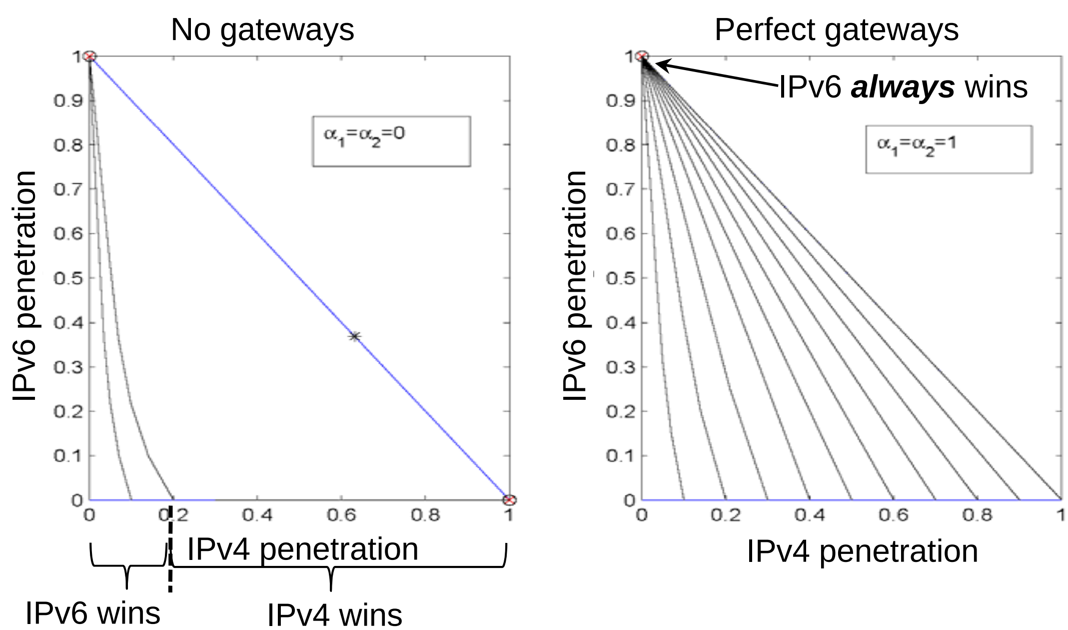

Gateways Can Also Help the Incumbent

- No gateways: Tech. 2 wipes out Tech. 1
- Perfect gateways: Tech. 1 nearly wipes out Tech. 2 (cannot eliminate it entirely though)
  - Note the large β value. This is a recurring theme for bad/odd things to happen, *i.e.,* heterogeneity in externalities is needed

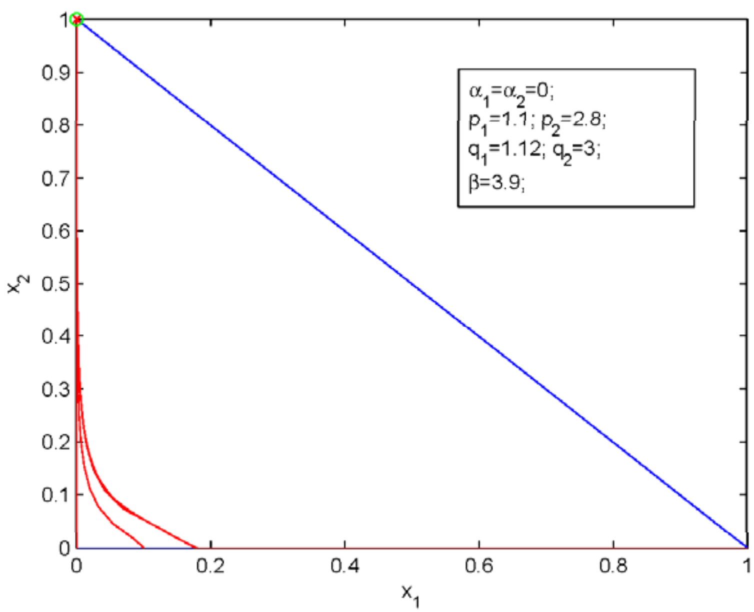

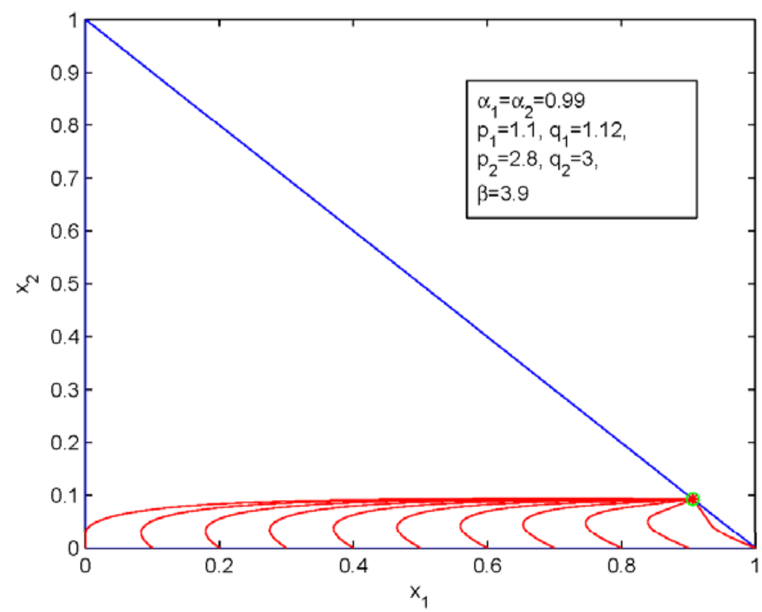

Key Findings – (1)

1. The system can have at most two stable equilibria (combinations of 2 out of 3 among Tech. 1 wins, Tech.2 wins, and Tech. 1 and Tech. 2 coexist)
   – When multiple equilibria are feasible, initial penetration determines the outcome (basins of attraction of each equilibrium)
2. Gateways can help either technology
   – Tech. 2 always benefits from better gateways, while better gateways can occasionally harm Tech. 1
3. Better gateways can harm overall penetration

Hurting Overall Market
(Asymmetric Gateways – Tech. 1)

- In the absence of gateways, Tech. 2 takes over the entire market
- Tech. 1 introduces gateways of increasing efficiency
  - Tech. 1 reemerges, but ultimately reduces overall market penetration

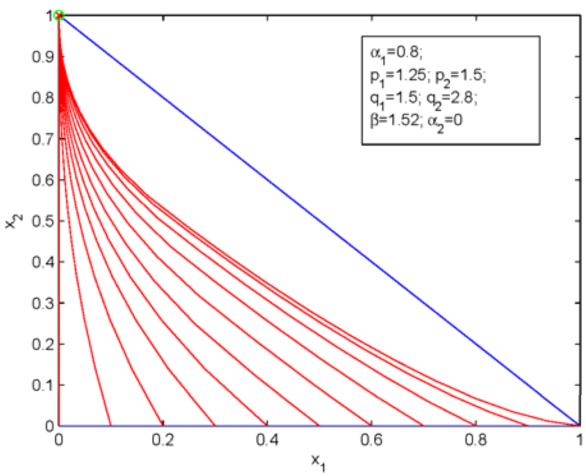

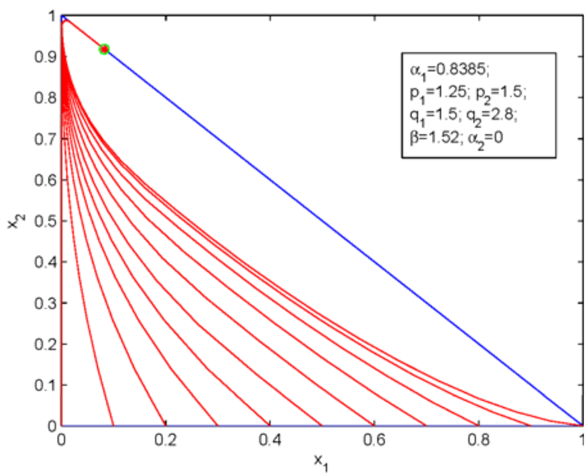

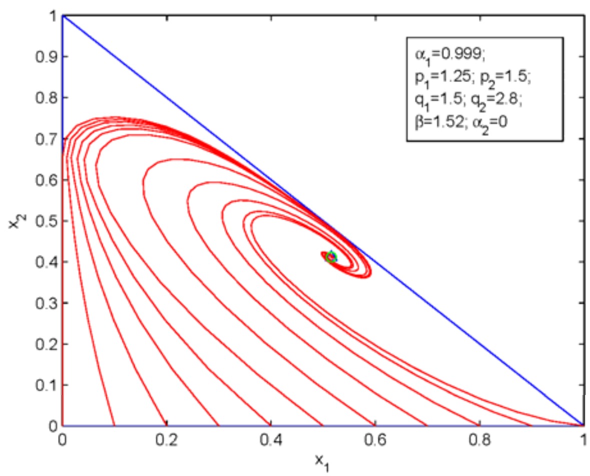

Hurting Overall Market
(Asymmetric Gateways – Tech. 2)

- Tech. 2 fails to gain market share without gateways
- Tech. 2 introduces gateways of increasing efficiency
  - Tech. 2 gains market share, but at the cost of a lower overall market penetration

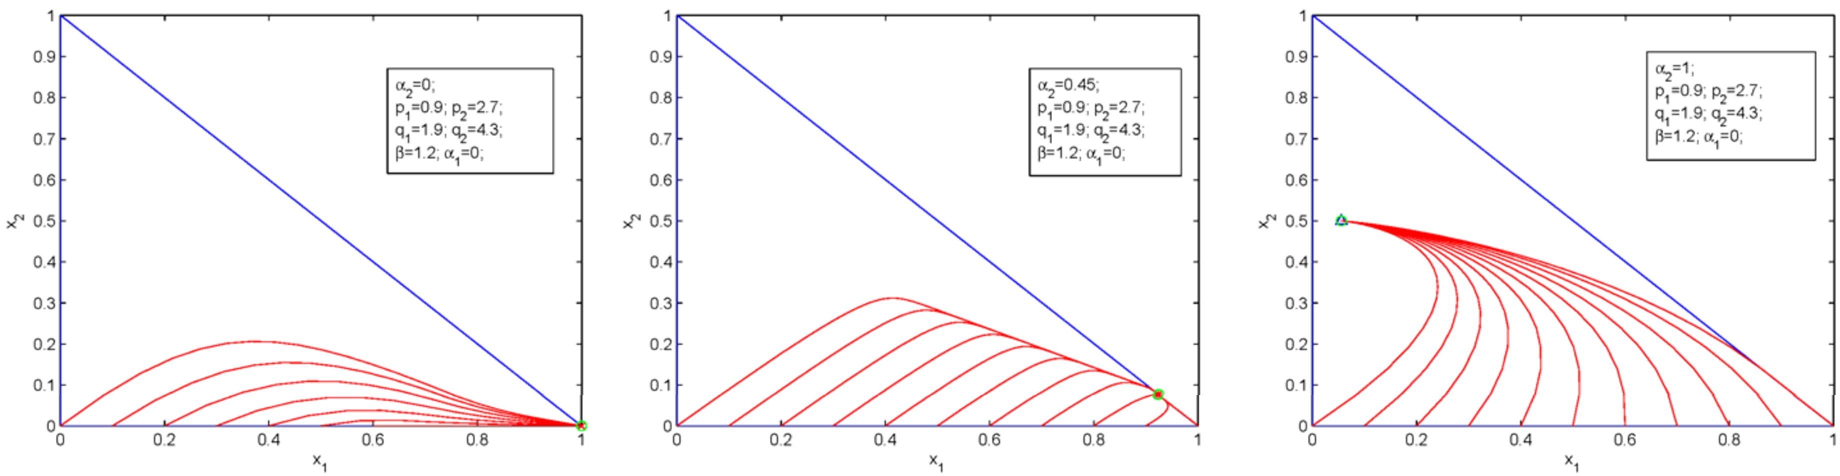

Hurting Overall Market
(Symmetric Gateways)

- Better gateways take Tech. 2
  - From 100% market penetration
  - To a combined market penetration below 20%!

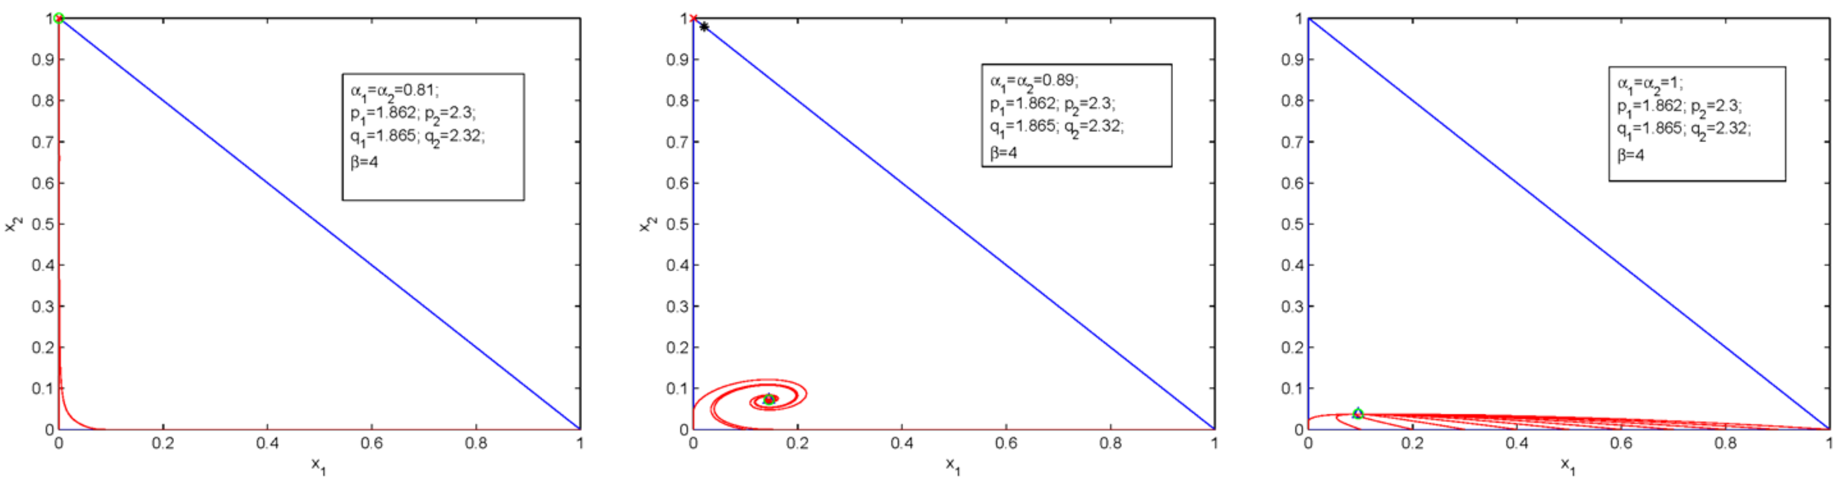

Key Findings – (2)

4. Gateways can prevent convergence of technology adoption (cyclical trajectories)
   - Does ***not*** arise without gateways
   - This requires heterogeneous technologies with $\alpha_1\beta > 1$, *i.e.,* Tech. 1 users derive higher externality benefits from Tech. 2 users than from other Tech. 1 users (the video-conf example)

Asymmetric Gateways

(From Stable to Unstable)

- As the efficiency of Tech. 1 gateway increases, system goes from dominance of Tech. 2 to a system with no stable state
  - No stable equilibrium for $\alpha_1 = 1$ and $\alpha_2 = 0$

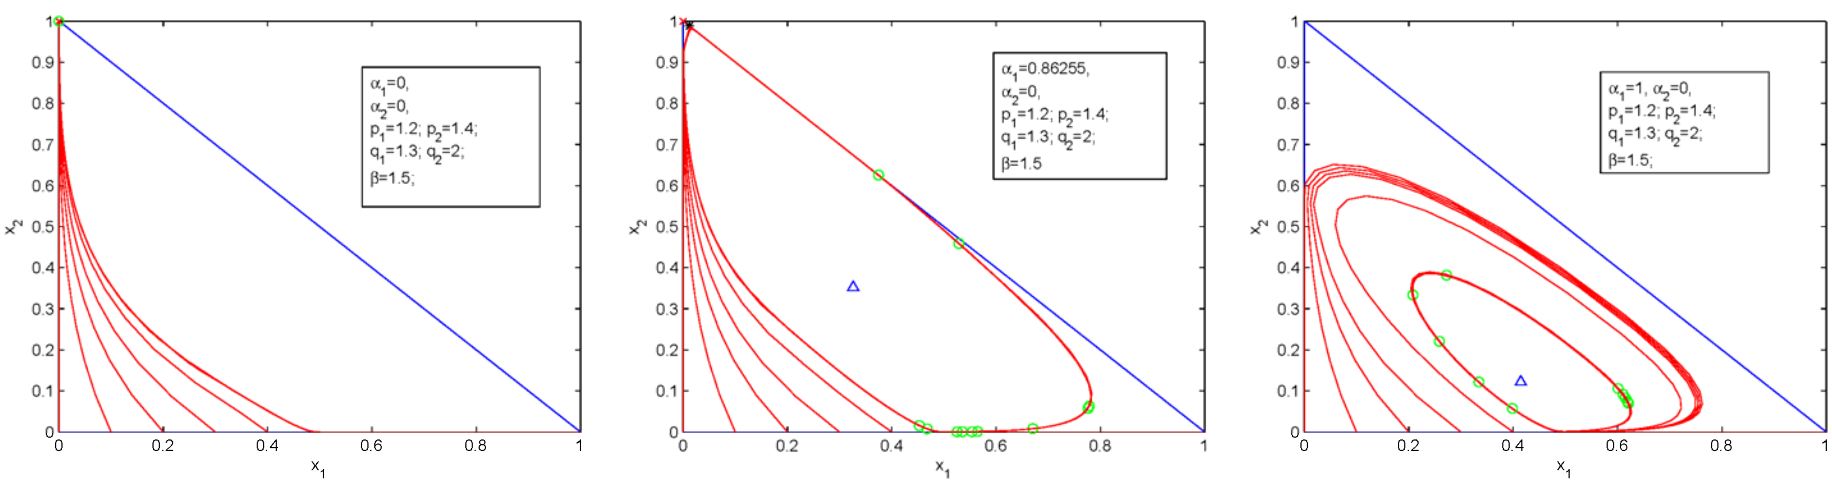

Symmetric Gateways

(From Stable to Unstable to Stable)

- No gateways: Tech. 2 captures full market
- Low efficiency gateways: No stable outcome
- Medium efficiency gateways: Neither tech. makes much inroad
- High efficiency gateways: Tech. 1 dominates at close to full market penetration

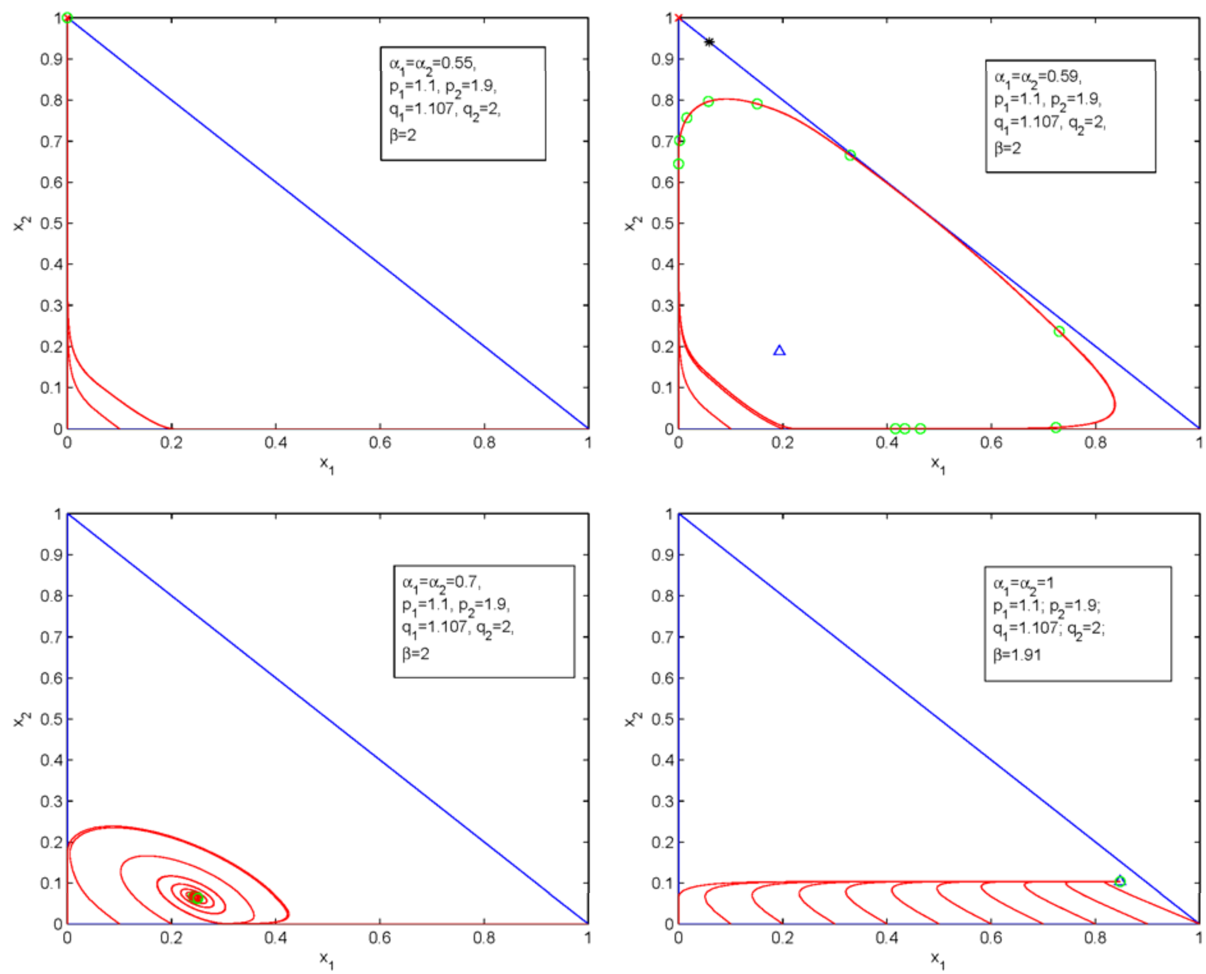

Results Robustness

- Most/all results hold for a wide range of model variations
  - No closed-form solutions, but numerical investigations are possible
- Model variations
  - Heterogeneity in user decisions (θ)
    - Non-uniform distributions
      - Positively and negatively skewed Beta-distributions
    - Extended to externality benefits
  - Other non-linear externality models
    - Sub- and super-linear: x^{α}, $0<\alpha<1$, and $\alpha>1$, logarithmic: $\log(x+1)$
    - Pure externalities (no intrinsic technology value)
  - Switching costs
    - Contract breaking penalties (present with any change)
    - Learning curve (incurred once)
    - Various combinations based on who and when they apply

WHAT IPv6 MIGRATION?

Why the Question (Now)?

- By any account we have had IPv6 for 15 years and have not bothered migrating to it, and for good reasons – there were really no incentives, but this is about to change
- By most accounts we are going to run out of IPv4 addresses soon (from http://www.potaroo.net/tools/ipv4 - and many others)
  - Projected IANA Unallocated Address Pool Exhaustion: **22-Sep-2011**
  - Projected RIR Unallocated Address Pool Exhaustion: **11-Jul-2012**

- This raises two important issues
  1. As IPv4 addresses become unavailable and we start using either IPv6 or private IPv4 addresses, we will need translation devices (NATs) to allow those hosts to reach the IPv4 Internet. **How big should those NATs be?**
  2. What incentives could we offer to foster migration of the IPv4 Internet to IPv6: **What about better network "quality"?**

How Much of the Internet Is Accessible Over IPv6 Today?

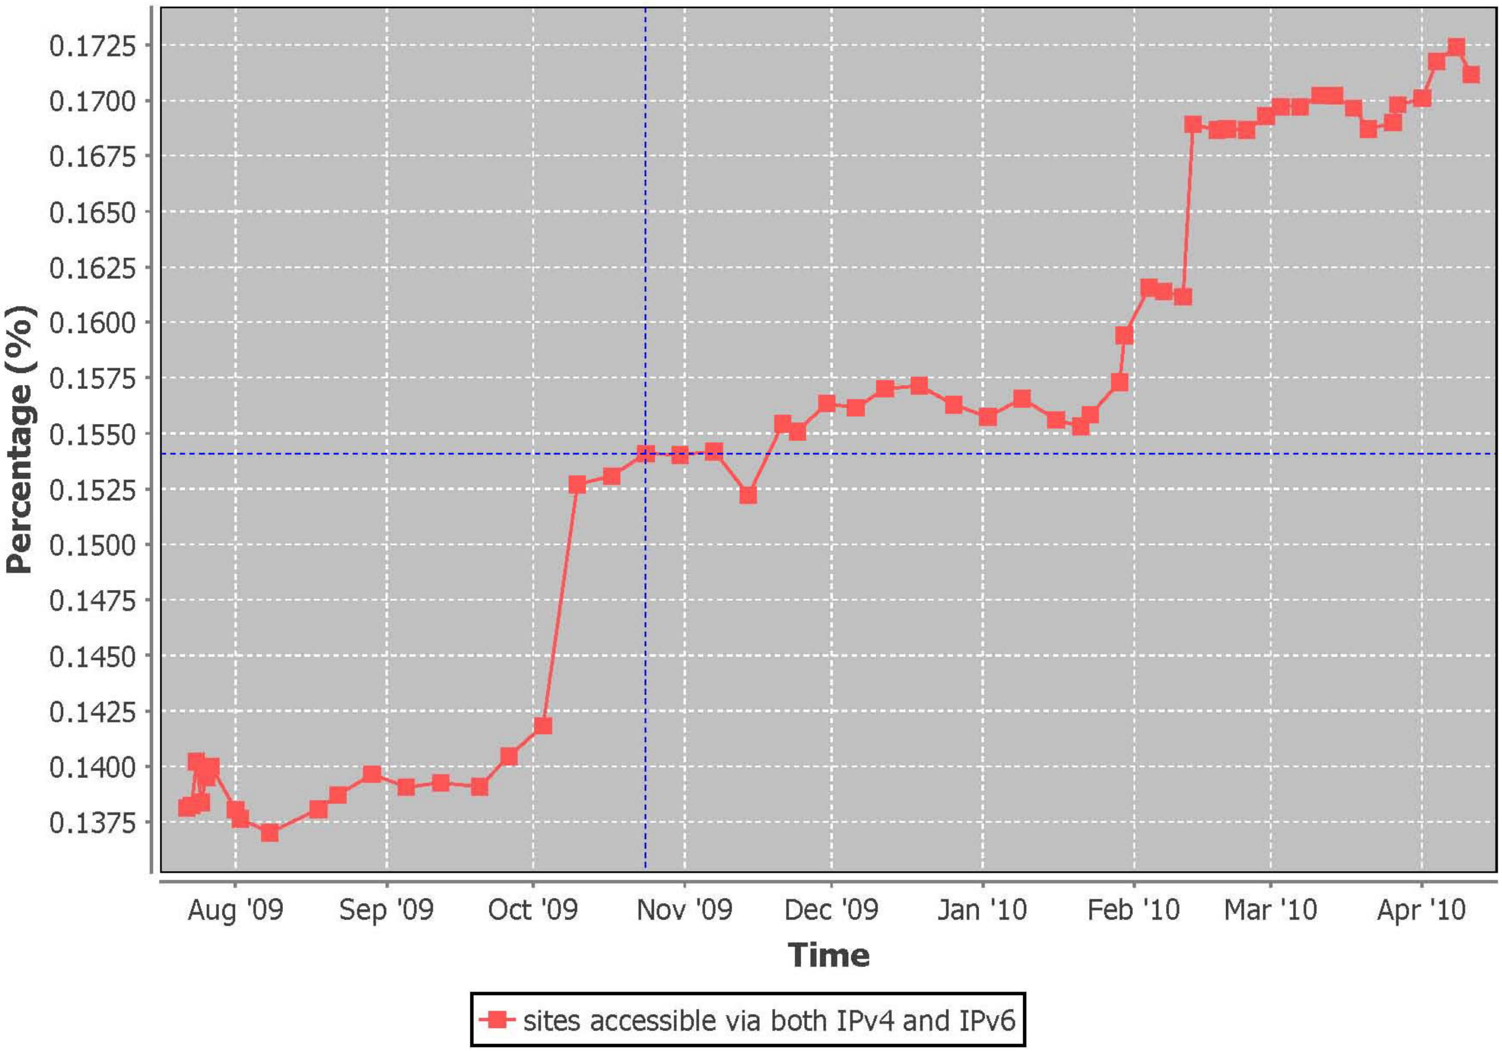

A Finer-Grain Look at IPv6 Reachability

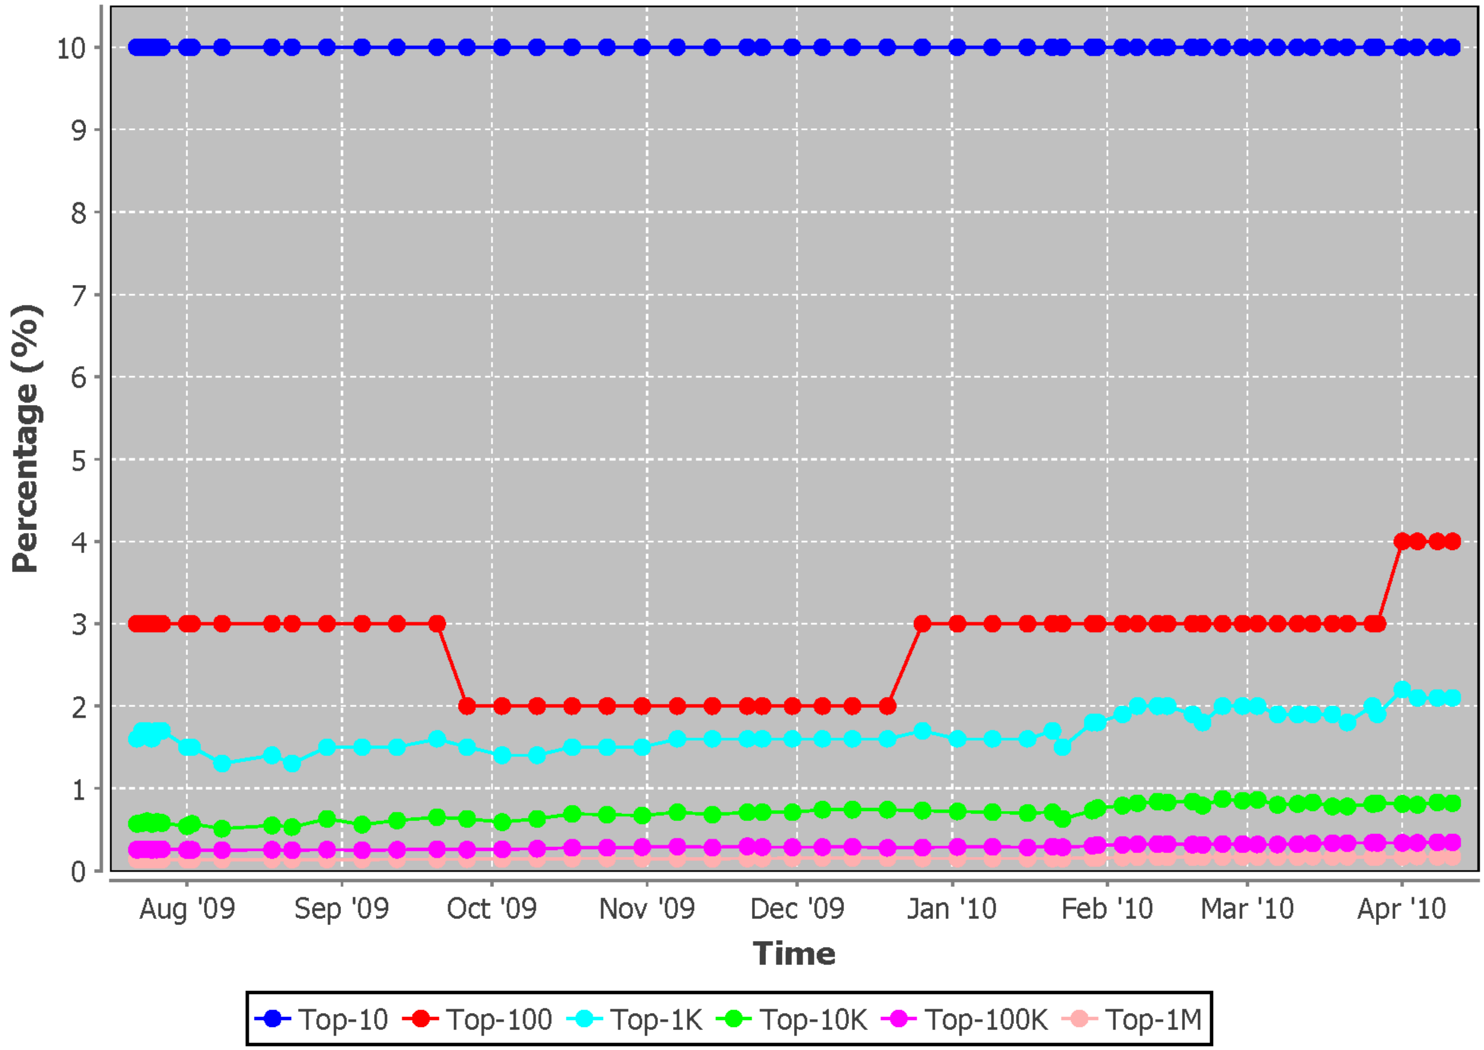

How Often Is IPv6 Better/Worse?

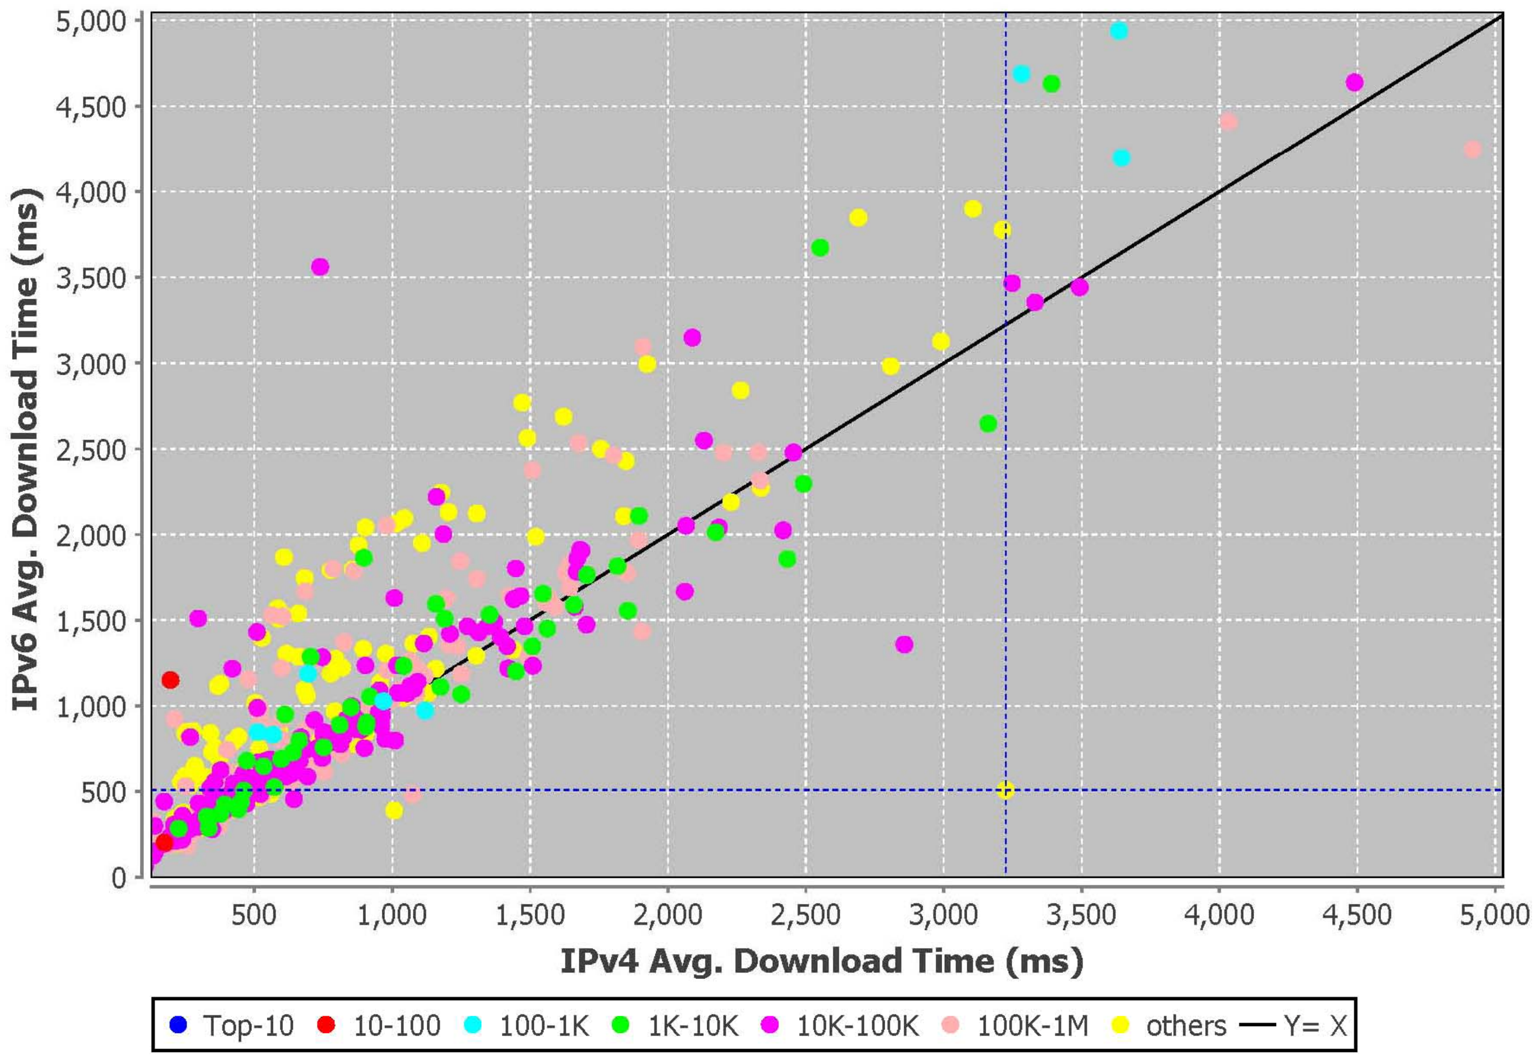

For Whom Is IPv6 Better?

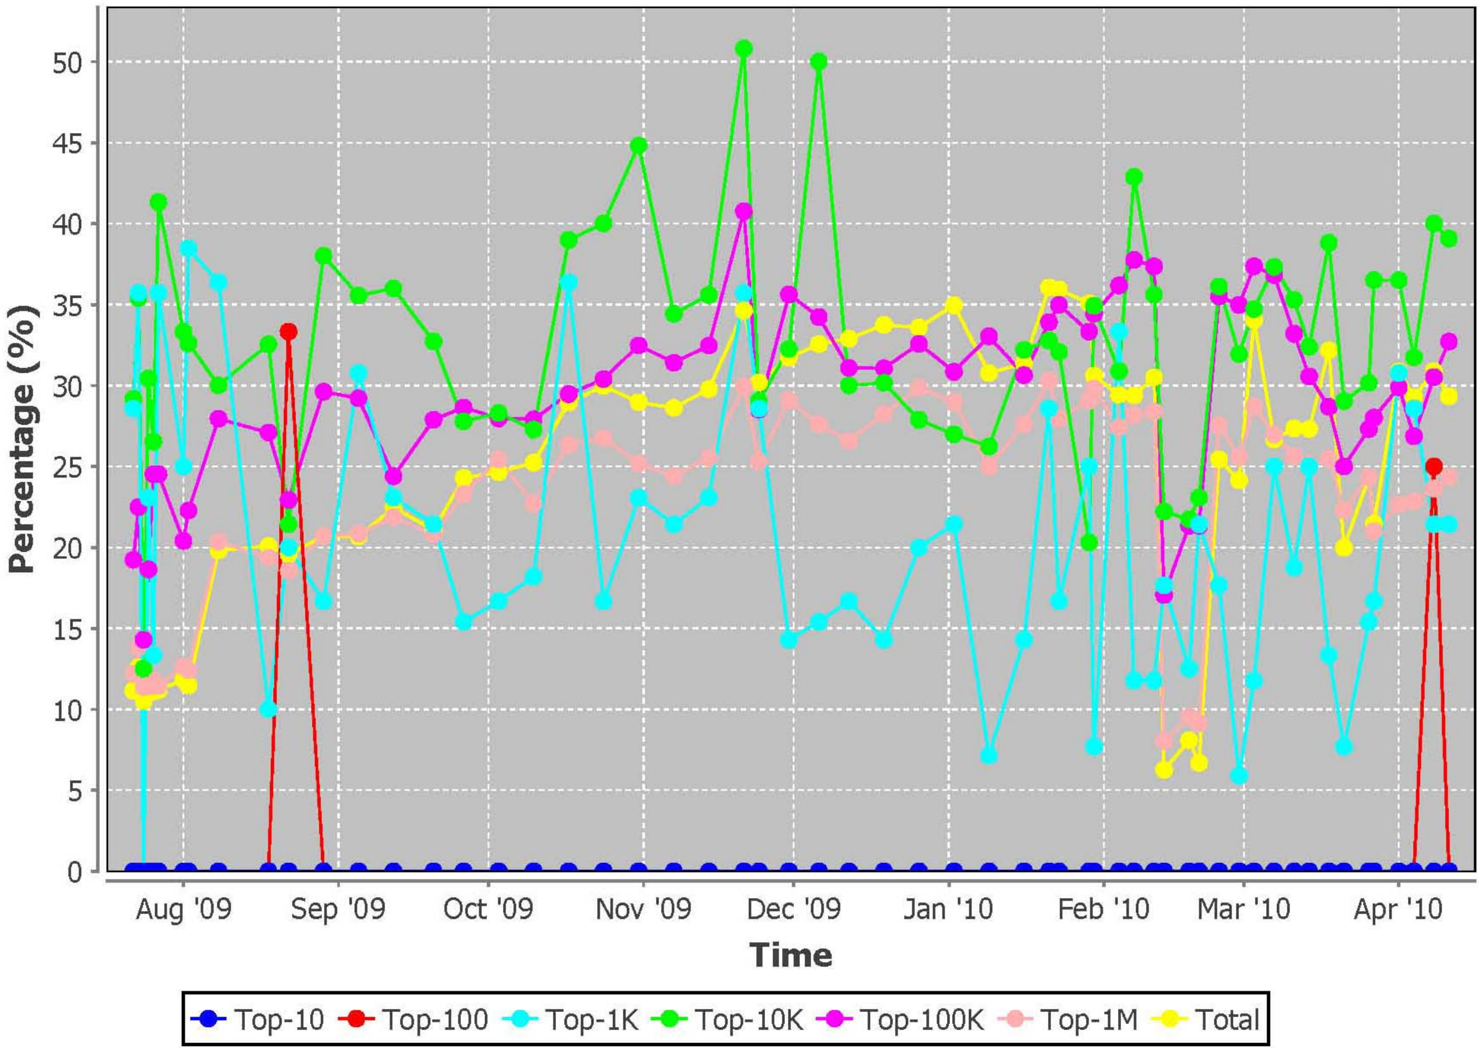

A Simplistic Model and Some Even More Simplistic Assumptions

- Three "players": Users, Internet Service Providers (ISPs), and Internet Content Providers (ICPs)
- Users rely on ISP(s) to access ICPs but are not decision makers
- ISPs have an existing base of IPv4 users, x_4, and a growing (exogenous) set of new IPv6 users, $x_6(t)$
  - ISP is responsible for provisioning translation capacity for IPv6 users: A ***decision variable*** a
  - ISP can also allocate IPv6 addresses to a fraction of existing IPv4 users: A ***decision variable*** $\alpha(t)$
- ICPs derive (advertising) revenues from users and incur connectivity configuration costs (variable across ICPs, *e.g.,* function of skill level)
  - Revenue proportional to number of users and "quality" of connectivity between them and ICP (q_{44}, q_{64}, q_{66}, where $q_{44}= 1$ and $q_{64} \leq q_{44}$) – homogenous across ICPs and users (***function of ISP(s) decisions***)
  - IPv6 connectivity incurs additional configuration cost ($\theta_i c_6$ for ICP i): A ***decision*** based on whether or not it improves revenue

ISP Cost – ICP Profit

- ISP translation cost (IPv6↔IPv4)

$$T(t) \cong x_6(t)(1-\beta(t))$$

  - $\beta(t)$ is fraction of ICPs that ***choose*** to be accessible over IPv6
- ICP Profit: Function of connectivity choices, *i.e.,* IPv4 only or IPv4+IPv6

$$\Pi_4^{(i)} = x_4 + q_{64}(t)x_6(t)$$

$$\Pi_6^{(i)} = x_4(1-\alpha(t)) + (\alpha(t)x_4 + x_6(t))q_{66}(t) - \theta_i c_6$$

- ICPs IPv6 connectivity decision

$$\beta(t) = \begin{cases} 0 & \text{if } q_{66}(t) < q_{64}(t) \le 1 \\ \left(\dfrac{\alpha(\mathrm{t})x_4(q_{66}(t)-1) + x_6(t)(q_{66}(t) - q_{64}(t))}{c_6} \right)_{[0,1]} & \text{otherwise} \end{cases}$$

Simple Findings from our Simple Model

- An obvious dilemma
  - We need good quality translation devices to make IPv6 acceptable to new users ($q_{64} \approx q_{44}$)
  - However, good translation quality makes it that much harder to convince ICPs to become IPv6 accessible (this requires $q_{66} > q_{64}$)

- A slightly less obvious question
  - Can we (ISPs) use connectivity quality differentials as a means of keeping the volume of translation traffic below the provisioned translation capacity a, and if yes, to what extent?

Controlling Translation Traffic Volume

- Keeping translation traffic bounded as the number of IPv6 users grows requires making native IPv6 connectivity ***better*** than what is achievable through translation devices

- How good native IPv6 connectivity needs to be depends on how bad translation quality is
  - If translation quality is close to that of the current IPv4 Internet (as it may have to), keeping translation traffic bounded calls for at some point making IPv6 connectivity ***better*** than that of IPv4
  - Once the number of IPv6 users is large enough, IPv6 connectivity quality can again be ***reduced*** down to that of IPv4 (or even lower…)

Impact of Translation Quality

Translation quality $q_{64} = 0.74$

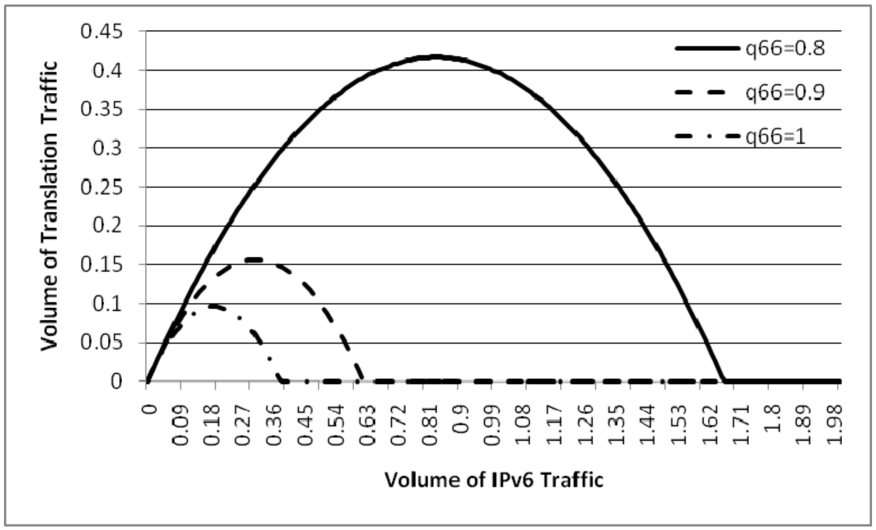

Translation quality $q_{64} = 0.79$

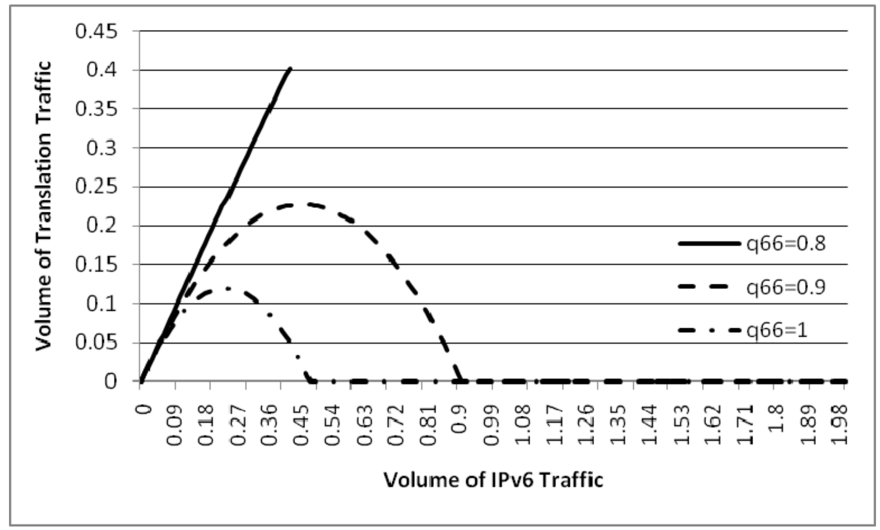

- As predicted, higher translation quality results in larger translation traffic volume or the need for better IPv6 connectivity

How Much Better Should IPv6 Be?

Goal is to keep translation traffic below 10% of current IPv4 traffic

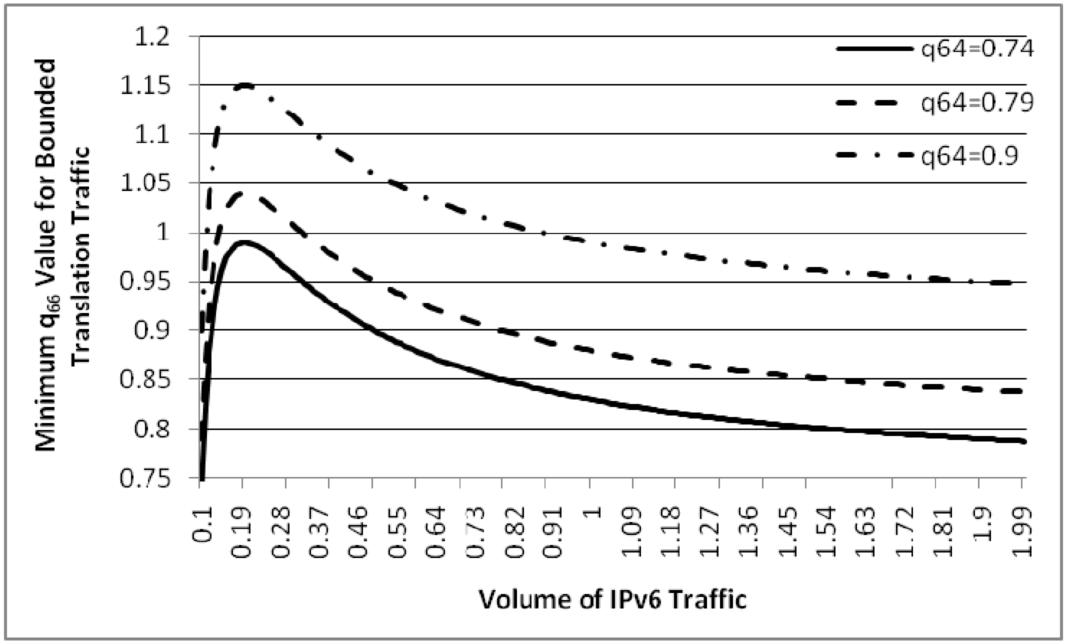

- Translation quality within 10% of native IPv4 access calls for making IPv6 connectivity about 15% better than IPv4 (when IPv6 users generate about 20% of the current IPv4 traffic)
- IPv6 quality can fall back to the quality level of the current Internet when the IPv6 user base is roughly the size of the current IPv4 Internet user base

Summary

- Adoption of (new) network technologies is no simple feat
  - Many interacting factors, several of which have nothing to do with technology itself
- Incompatible technologies that require gateways make matters even worse
  - Gateways can induce unexpected/undesirable effects
  - Gateways are no panacea – They need to be good initially, but if they are too good, migration may never happen unless the new technology offers a clear advantage
    - The IPv6 scenario provides a good example, where unless we are able to offer a real incentive for current IPv4 users (ICPs) to migrate, *e.g.,* in the form of better quality, they may never do!