6-4-2010

Fostering IPv6 Migration Through Network Quality Differentials

Roch A. Guérin
University of Pennsylvania, guerin@acm.org

Kartik Hosanagar
University of Pennsylvania, kartikh@wharton.upenn.edu

Recommended Citation

Fostering IPv6 Migration Through Network Quality Differentials

Abstract

Although IPv6 has been the next generation Internet protocol for nearly 15 years, new evidences indicate that transitioning from IPv4 to IPv6 is about to become a more pressing issue. This paper attempts to quantify if and how such a transition may unfold. The focus is on ``connectivity quality,'' e.g., as measured by users' experience when accessing content, as a possible incentive (or disincentive) for migrating to IPv6, and on ``translation costs'' (between IPv6 and IPv4) that Internet Service Providers will incur during this transition. The paper develops a simple model that captures some of the underlying interactions, and highlights the ambiguous role of translation gateways that can either help or discourage IPv6 adoption. The paper is an initial foray in the complex and often puzzling issue of migrating the current Internet to a new version with which it is incompatible.

Keywords

IPv6, migration, incentives, network quality

Disciplines

Digital Communications and Networking | Management Information Systems | OS and Networks | Other Economics | Technology and Innovation

Comments

Fostering IPv6 Migration Through Network Quality Differentials*

Roch Guérin
Elec. & Sys. Eng, U. Pennsylvania
guerin@ee.upenn.edu

Kartik Hosanagar
Wharton, U. Pennsylvania
kartikh@wharton.upenn.edu

ABSTRACT

Although IPv6 has been the next generation Internet protocol for nearly 15 years, new evidences indicate that transitioning from IPv4 to IPv6 is about to become a more pressing issue. This paper attempts to quantify if and how such a transition may unfold. The focus is on "connectivity quality," *e.g.,* as measured by users' experience when accessing content, as a possible incentive (or disincentive) for migrating to IPv6, and on "translation costs" (between IPv6 and IPv4) that Internet Service Providers will incur during this transition. The paper develops a simple model that captures some of the underlying interactions, and highlights the ambiguous role of translation gateways that can either help or discourage IPv6 adoption. The paper is an initial foray in the complex and often puzzling issue of migrating the current Internet to a new version with which it is incompatible.

Categories and Subject Descriptors

C.2.3 [**Network Operation**]: Public networks; C.2.6 [**Internetworking**]: Standards

General Terms

Economics, Management, Standardization

Keywords

IPv6, migration, incentives, quality

1. BACKGROUND AND MOTIVATIONS

IPv6, the next generation Internet Protocol, was standardized about fifteen years ago [4] to address a number of deficiencies in the current (IPv4) version of the protocol. The changes included, among others, the addition of built-in security capabilities, improved support for mobility, a more streamlined handling of options, simpler mechanisms for address allocation, and a larger address space (128 bits vs. 32 bits). However, the limited adoption of IPv6 to-date[1] indicates that these various enhancements have not been a sufficient incentive to trigger widespread adoption of IPv6. This is in part because IPv4 has itself been extended to include or at least approximate many of these enhancements, *e.g.,* IPSEC offers comparable security capabilities in IPv4. As a result, the primary remaining differentiator of IPv6 is its larger address space, which has so far not proved sufficient to justify its adoption. This may be about to change. In particular, while a total of $2^{32} \approx 4.3$ billions IPv4 addresses may seem plentiful, we are fast approaching this limit, *e.g.,* http:/www.potaroo.net/tools/ipv4 has the Internet Assigned Numbers Authority (IANA) exhausting its address pool by September 2011, and Regional Internet Registries (RIRs) exhausting theirs about a year later.

Hence, as we enter a new era of IPv4 address scarcity, allocating IPv6 addresses is an option that, if not yet desirable, is becoming viable (alternatives such as private addresses have, among other disadvantages, translation costs that keep rising with usage, so that they don't offer a true long-term solution). Translating this option into reality is, however, no simple feat given its limited success to date (see [8] for an insightful discussion on this issue and [7] for a review of various alternatives). Specifically, making IPv6 a reality involves several considerations.

First, IPv6 needs to be supported across the many devices that connect to the Internet and are used to build it. This is by now largely a reality, thanks in part to a June 30, 2008, US Government's Office of Management and Budget mandate for federal agencies to be running IPv6 [3], and similar initiatives in other countries as reported in [10]. As a result, IPv6 support is now standard in pretty much all networking equipment and major operating systems[2], as well as in numerous applications and consumer devices (see http://www.ipv6-to-standard.org for a reasonably comprehensive list).

The next key component in realizing an IPv6 Internet is the network itself. This goes beyond equipment capabilities and touches on many issues related to configuration, management and policies. The Domain Name System (DNS) itself is by now largely IPv6 capable even if a number of local servers have not yet been upgraded. On the other hand, the network infrastructure, and in particular routing and peering agreements, is less advanced in spite of a growing number of Autonomous Systems (ASes) experimenting with IPv6 (see the RIPE Labs IPv6 measurement page at http://labs.ripe.net/content/ipv6-measurement-compilation for a list of measurement efforts assessing IPv6 deployments).

Finally and possibly most importantly, an IPv6 Internet calls for unimpeded access to all the existing IPv4 *content,* as content is arguably a major component of what defines the value of the Internet. This consideration stems from the

---

[1]See [9] or http://mnlab-ipv6.seas.upenn.edu/monitor for a snapshot of Internet content accessibility over IPv6.

[2]Unix, Linux, and Mac OS X have supported IPv6 for many years, and the Vista release has finally brought solid IPv6 support to Windows.

incompatibility of IPv4 and IPv6, which results in IPv6-only devices, *i.e.,* devices without an IPv4 address, being unable to directly access devices or content reachable only over IPv4. Specifically, an IPv6-only device seeking to connect to a "site" queries DNS for the IP address associated with the site's name, but does so requesting an address of type AAAA (quad-A). In the absence of a registered IPv6 address for the site, DNS will inform the user that the site is unreachable (over IPv6). IPv6 connectivity is, therefore, of little or no value to IPv6 users without gateways or "translation" devices, *e.g.,* [6], that provide access to the IPv4 Internet. Furthermore, poor IPv6 accessibility of content also affects the capacity requirements and, therefore, costs of those gateways. This is because the volume of traffic requiring translation grows with both the number of IPv6-only users and the amount of content (number of sites) inaccessible over IPv6 (see Section 2.2 for details).

Fostering IPv6 accessibility among Internet content providers is, therefore, vital to an eventual migration to an IPv6 Internet. However, there are few if any intrinsic incentives that can motivate current content providers to make themselves accessible over IPv6[3]. Furthermore, content providers are unlikely to even consider IPv6 unless it is clear that it will not affect how existing (IPv4) users are able to access them. This is well illustrated by the procedures that Google has recently put in place for IPv6 access to its services. Specifically, Google has made itself accessible over IPv6, but through a different (and more limited) service at http://ipv6.google.com instead of the standard http://www.google com, *except* for users connected to a service provider certified by Google as having *good* IPv6 connectivity (see http://www.google.com/intl/en/ipv6/ for details). An exact definition of good connectivity quality is likely to be elusive, but the "Google over IPv6" page provides a useful summary of the main criteria, *i.e.,* "Low-latency, redundant paths using direct peering or reliable transit." In other words, the experience of a user accessing a web site over IPv6 should be comparable to that of a user accessing it over IPv4. Unfortunately, preliminary measurements[4] seem to indicate that this is often not the case, with roughly 75% of the web sites accessible over both IPv4 and IPv6 having worse page download times over IPv6 than over IPv4. This is hardly an incentive for content providers to consider making themselves accessible over IPv6.

In the rest of this paper, we develop a simple model to investigate how connectivity quality, *e.g.,* content download speed and reliability, may affect the eventual adoption of IPv6. The results from the model help elucidate how connectivity quality can impact both IPv6 adoption by content providers, and the volume of IPv6↔IPv4 translation traffic that service providers will need to handle while transitioning from an IPv4 Internet to an IPv6 one.

The rest of the paper is structured as follows. Section 2 introduces our simple model, its notations and the assumptions on which it relies. The model is analyzed and discussed in Section 3, which illustrates how connectivity quality affects both IPv6 adoption and the volume of translation traffic. Section 4 summarizes our main findings and their implications for IPv6 adoption. It also points to possible

---

[3]They already have an IPv4 address, so that the larger IPv6 address space is of little significance to them.

[4]See http://mnlab-ipv6.seas.upenn.edu or
http://ipv6monitor.comcast.net.

extensions to generalize the model and/or eliminate some of the simplifying assumptions on which it relies.

2. MODEL AND NOTATION

Because of our focus on IPv6 adoption by Internet Content Providers (ICPs) and on the provisioning requirements of the IPv6↔IPv4 gateways that Internet Service Providers (ISPs) need to deploy, ISPs and ICPs are the decision makers in the model. Users are present but only as an exogenous parameter that may affect ISP and ICP decisions. In other words, users care about their ability to access Internet content, but are mostly oblivious to how this is realized, *i.e.,* over IPv6 or IPv4. In particular, users access ICPs that derive (advertising) revenues from them, and growth in the number of new users to which an ISP needs to allocate IPv6 addresses (because of a shortage of IPv4 addresses) partly affects the volume of translation traffic that gateways handle. The next subsection discusses in more details interactions between ISPs and ICPs and the role of users, while Subsection 2.2 introduces the model itself.

2.1 Assumptions and Notation

We assume an ISP (more generally a set of ISPs) that because of shortage of IPv4 addresses has started allocating IPv6 addresses to new users. Implicit in this choice is a preference for IPv6 over private IPv4 addresses. As mentioned earlier, this is because private addresses incur similar "translation costs" as IPv6 addresses when connecting to the public IPv4 Internet, *i.e.,* from a private IPv4 address to a public one and back, without the possibility of an eventual migration to a translation-free environment. In other words, while private addresses may offer short-term benefits, *i.e.,* the use of a familiar technology, their long-term costs keep growing with the size of the Internet. In contrast, IPv6 incurs short-term deployment and training costs, but has the potential for much lower long-term costs. Hence, it can be argued to be the better option if its long-term benefits can be realized. Investigating how this can be accomplished is one of the paper's motivations. Obviously, in this model the ISP is assumed to operate a network that is both IPv4 and IPv6 capable, and therefore to have made the necessary equipment and operational investments.

The ISP has two categories of users: Existing users that have been allocated an IPv4 address (and possibly also an IPv6 address), and new users that only have an IPv6 address. As mentioned above, users have no control over (and no interest in) the type of IP address they receive and are not decision makers in the model. The size of the IPv4 and IPv6 user populations are denoted as x_4 and $x_6(t)$, respectively, where x_4 is fixed (the user population when the ISP ran out of IPv4 addresses), and $x_6(t)$ is an exogenous, non-decreasing function of time (t) with $x_6(0) = 0$, *i.e.,* user demand is unaffected by ISPs and ICPs decisions. In addition to new users that have an IPv6 address, a fraction $\alpha(t), 0 \leq \alpha(t) \leq 1$, of the existing x_4 IPv4 users is assumed to be also IPv6 accessible, *i.e.,* have been assigned both IPv4 and IPv6 addresses by the ISP (as mentioned earlier, most access devices nowadays support both IPv4 and IPv6). $\alpha(t)$ is a decision variable of the ISP.

Users are assumed homogeneous in how they access ICPs and in the amount of traffic they generate. ICPs derive (advertising) revenues from users accessing them, and are also assumed homogeneous in their revenues and in the amount

of user traffic they sink and source. This is obviously an over-simplification as ICPs vary in popularity, *e.g.,* [1], which influences traffic volume, but should not significantly affect conclusions. IPv6 adoption by a popular site could be accounted for in our simple model as adoption by multiple sites. Transition points may shift, but general behaviors should remain similar.

All ICPs own an IPv4 address (registered with DNS), and can decide to also register an IPv6 address to enable native IPv6 access[5]. Implicit here is the assumption that IPv6 connectivity is ubiquitous. This is not true today, as many ASes are not yet IPv6 reachable. However, many of the largest ISPs are IPv6 enabled (see http://bgp.he.net/ipv6-progress-report.cgi), and connectivity to IPv6 enabled sites is broadly available, *e.g.,* in over 6.5 million tests, only 0.13% identified problems for users trying to reach a web site after it became IPv6 reachable (see http://ipv6test.max.nl).

An ICP's decision on whether or not to become IPv6 accessible depends on two factors. The first is the cost of adding IPv6 access. In the model, this cost is denoted as $\theta_i c_6$, where c_6 denotes a base cost, i is an index that identifies ICP i, and the variable θ_i accounts for heterogeneity across ICPs. The second factor is the impact of becoming IPv6 accessible on the ICP's revenue. This is where connectivity quality is taken into account. Specifically, ICP revenue R is assumed to be of the form $R \sim qx$, where x denotes the number of users accessing the ICP and q is the quality of their connectivity to the ISP. Revenues are increasing in connectivity quality to reflect that users may spend more time browsing content as network quality increases and therefore generate higher (advertising) revenues for the ICP. In traditional econ models, this captures the effect of the Internet experience becoming more attractive than the "outside good" (*e.g.,* TV and other substitutes). ICPs decide to become IPv6 accessible if the associated increase in revenues exceeds the cost.

With regard to network quality, it is necessary to distinguish between three possible types of network connectivity and associated quality levels: q_{44}, $q_{64}(t)$ and $q_{66}(t)$. They correspond to the three possible connectivity combinations between users and ICPs, namely, IPv4↔IPv4, IPv6↔IPv4, IPv6↔IPv6, respectively.

Connectivity quality of the IPv4 Internet is assumed fixed, *i.e.,* $q_{44} = 1$, while connectivity quality through translation gateways, $q_{64}(t)$, and native IPv6 connectivity, $q_{66}(t)$, can vary, *e.g.,* $q_{64}(t)$ can decrease as translation boxes become more heavily loaded and conversely $q_{66}(t)$ can improve as technology and skills mature. We assume that $q_{64}(t) \leq q_{44}$, *i.e.,* translation can only lower the quality of accessing content over the IPv4 Internet. However, there is no such constraint on the quality of IPv6 connectivity, $q_{66}(t)$, that can be better or worse than IPv4 connectivity with or without translation. For example, coarser peering agreements could force IPv6 traffic onto longer paths resulting in longer delays, *i.e.,* $q_{66}(t) < q_{44}$. Alternatively, the small number of IPv6 users or the decision by the ISP to give precedence to IPv6 traffic could initially ensure that IPv6 traffic sees lower congestion levels than IPv4 traffic, which would in turn result in $q_{66}(t) > q_{44}$. We assume that the ISP, or more generally ISPs, can control the value of $q_{66}(t)$, *e.g.,* through

provisioning and configuration of their equipment[6], and by their choices of peering decisions. Similarly, ISPs also control the value of $q_{64}(t)$, *i.e.,* translation quality; for example by provisioning translation capacity or more generally by requiring that equipment vendors meet specific performance benchmarks.

We note that the notion of global connectivity quality metrics such as q_{44}, $q_{64}(t)$, and $q_{66}(t)$ is obviously a simplification. Connectivity quality is an end-to-end property that varies across users and sites. The quantities q_{44}, $q_{64}(t)$, and $q_{66}(t)$ are instead meant to capture average metrics, for which relative comparisons are what matters, *i.e.,* to what extent does connectivity quality in the current Internet differ from what is achievable when crossing translation gateways and/or relying on native IPv6 connectivity.

In the next section, we formalize the above discussion by introducing a simple model that captures the decision process of both ICPs and ISPs.

2.2 Model Formulation

2.2.1 ISP Decision Process and Variables

As stated in the previous section, we consider an ISP that has already made the decision to adopt IPv6, *e.g.,* to be able to allocate IPv6 addresses to new users. The ISP still needs to decide how to set $q_{64}(t)$ and $q_{66}(t)$. As discussed above, different values for $q_{64}(t)$ and $q_{66}(t)$ can be realized based on how equipment is provisioned and configured, and through decisions on how IPv6 traffic is to be routed. These choices have costs of their own and also affect translation costs $T(t)$, *i.e.,* the costs of the translation gateways required to allow new IPv6 users to access the legacy IPv4 Internet. Those translation costs grow with the volume of translation traffic, and are assumed to be of the form:

$$T(t) \sim x_6(t)\,(1 - \beta(t))\ , \tag{1}$$

where the parameter $\beta(t)$ denotes the fraction of ICPs that have decided to become IPv6 accessible by time t. In particular, Eq. (1) reflects the fact that if all content was accessible over IPv6 ($\beta(t) = 1$), there would be no need for translation; the original assumption when IPv6 was first standardized.

Additionally, the ISPs may also decide to allocate IPv6 addresses to existing IPv4 users, as captured through the variable $\alpha(t)$. The ISPs main concern in making such a choice as well as in choosing $q_{64}(t)$ and $q_{66}(t)$ is the long-term impact on $T(t)$, and in particular the extent to which it is possible to keep it bounded as $x_6(t)$ increases. Note that this requires that $\beta(t)$ itself increase, *i.e.,* as the number of IPv6 users grows, more of the current Internet content needs to become IPv6 accessible. Understanding how $\alpha(t), q_{64}(t)$ and $q_{66}(t)$ influence $\beta(t)$ calls for modeling the decision process of ICPs. This is the topic of the next section.

2.2.2 ICP Decision Process and Variables

As discussed earlier, ICPs derive revenues that are proportional to the quality of user access, *i.e.,* $q_{44}, q_{64}(t)$ and $q_{66}(t)$. They also incur costs when deciding to become IPv6 accessible. ICPs, therefore, evaluate connectivity options (IPv4 only or both IPv4 and IPv6) and select the one with the highest profit. When all users have either an IPv4 or an IPv6 address but not both, this evaluation is straightforward. However, when some users have both IPv4 and

---

[5]For simplicity, the model ignore IPv6-only ICPs (their number is anyhow marginal).

[6]What share of bandwidth or what priority level.

IPv6 addresses (the fraction $\alpha(t)$ of the previous section), it becomes necessary to specify what connectivity option they choose when a choice is available, *i.e.,* when accessing an ICP with both IPv4 and IPv6 addresses.

One model (model 1) assumes that IPv6 connectivity has precedence over IPv4. This is consistent with the default policy of [5]. In other words, when faced with a choice, end-systems are by default often configured to first try to connect using IPv6. Given that users are unlikely to bother changing default configurations, this is likely to be a common scenario. Another model (model 2) lets users choose the connectivity option with the higher quality. This corresponds to a rational decision process by users or providers of end-user devices based on awareness of quality differentials across connectivity options. For example, if IPv6 quality is well-known to be poor, providers of operating systems may begin shipping them configured to always prefer IPv4 whenever there is a choice, or conversely system administrators may choose this as the default configuration for systems they manage. Alternatively, end-system solutions may be developed that will allow applications to always select the best performing connectivity option (see [2] for a very insightful presentation on such an approach implemented in Apple's products). In spite of their apparent differences, the two models have mostly parallel analyses and broadly yield similar conclusions. As a result, we only present model 1 and its analysis, and point out differences between the two models when stating results.

The next two equations provide expressions for an ICP's profit as a function of whether or not it decides to become IPv6 accessible.

$$\Pi_4^{(i)} = x_4 + q_{64}(t)x_6(t) \qquad (2)$$
$$\Pi_{46}^{(i)} = x_4\left(1 - \alpha(t)\right) \qquad (3)$$
$$+ \ \left(\alpha(t)x_4 + x_6(t)\right)q_{66}(t) - \theta_i c_6$$

Eq. (2) is the ICP's profit if it remains only IPv4 accessible, while Eq. (3) gives its profit once it also becomes IPv6 accessible. Eq. (3) reflects the precedence of IPv6 over IPv4 of model 1.

ICP i decides to become IPv6 accessible only if it yields a higher profit, *i.e.,* $\Pi_{46}^{(i)} > \Pi_4^{(i)}$. Substituting the expressions for $\Pi_{46}^{(i)}$ and $\Pi_{46}^{(i)}$ of Eqs. (2) and (3) and assuming for simplicity of analysis that θ_i is uniformly distributed in $[0,1]$, an ICP becomes IPv6 accessible if $\theta_i \leq \theta^*$, where

$$\theta^* = \frac{\alpha(t)x_4(q_{66}(t) - 1) + x_6(t)(q_{66}(t) - q_{64}(t))}{c_6} \qquad (4)$$

3. ANALYSIS AND DISCUSSION

3.1 Influencing IPv6 Adoption

From Eq. (4), we get the following expression for the level $\beta(t)$ of IPv6 adoption by ICPs

$$\beta(t) = \left(\frac{\alpha(t)x_4(q_{66}(t) - 1) + x_6(t)(q_{66}(t) - q_{64}(t))}{c_6}\right)_{[0,1]}, \quad (5)$$

where we have used the notation $(x)_{[0,1]}$ to indicate the projection of x on the interval $[0,1]$.

Eq. (5) provides intuitive confirmation of the impact of the different parameters on IPv6 adoption by ICPs.

The first intuitive finding from Eq. (5), is that unless native IPv6 connectivity is better than what is achievable through translation devices, *i.e.,* $q_{66}(t) > q_{64}(t)$, current ICPs will never have any incentive to become IPv6 accessible. This presents a dilemma for ISPs, which may struggle with their early IPv6 deployment, *i.e.,* $q_{66}(t)$ is likely to initially be less than $q_{44} = 1$, while translation quality ($q_{64}(t)$) may at first be relatively high as translation traffic volume will be low. In short, ISPs need to be aware that they wont entice ICPs to start adopting IPv6 until their IPv6 infrastructure offers better connectivity than what is achievable through translation gateways, which will typically themselves be required to deliver quality close to that of the current Internet, *i.e.,* $q_{64}(t) \approx q_{44} = 1$.

Conversely, when IPv6 connectivity becomes better than what is available over the current IPv4 Internet, *i.e.,* $q_{66}(t) > q_{44} = 1$, then this alone is enough incentive for some ICPs to consider adopting IPv6 even when the number of IPv6 users is small. In other words, offering better connectivity over IPv6 than IPv4 is an effective tool for bootstrapping IPv6 adoption by current ICPs. This effect is reinforced if the ISP also starts allocating IPv6 addresses to existing IPv4 users, *i.e.,* $\alpha(t) > 0$. This is because it inflates the set of (IPv6) users that can benefit from the higher quality connectivity of IPv6; hence increasing the revenue gains that an ICP can realize by becoming IPv6 accessible.

On the other hand, this effect is reversed when IPv6 quality remains below that of the current IPv4 Internet, *i.e.,* $q_{66}(t) \leq q_{44} = 1$. In this case, providing existing IPv4 users with IPv6 addresses results in lower IPv6 adoption (smaller β values) by ICPs. This is due to the assumption that users with both IPv4 and IPv6 addresses give precedence to IPv6 access when that option is available (model 1). This forces users with IPv4 and IPv6 addresses to connect to IPv6 accessible sites using IPv6, and in the process experience lower quality than if they had used IPv4. This represents a disincentive for ICPs to become IPv6 accessible, even when IPv6 offers IPv6-only users better connectivity quality than translation devices, *i.e.,* $q_{66}(t) > q_{64}(t)$. In this scenario, providing better quality connectivity to IPv6 users is insufficient incentive to compensate for the poorer quality it imposes on users that have both IPv4 and IPv6 addresses. Specifically, when $q_{66}(t) < 1$, if the number $\alpha(t)$ of users with both IPv4 and IPv6 addresses is such that

$$\alpha(t) \geq \frac{x_6(t)\left(q_{66}(t) - q_{64}(t)\right)}{x_4\left(1 - q_{66}(t)\right)},$$

the lower quality these users experience with ICPs that are IPv6 accessible, is a sufficient disincentive to prevent all ICPs from becoming IPv6 accessible, *i.e.,* $\beta(t) = 0$.

As a result, the best strategy to maximize IPv6 adoption is to avoid enabling IPv6 access for existing IPv4 users, *i.e.,* keep $\alpha(t) = 0$, as long as $q_{66} \leq 1$. Note that this constraint is absent under the connectivity assumption of model 2, *i.e.,* users always select the better connectivity option, where setting $\alpha(t) > 0$ does not negatively affect IPv6 adoption even when $q_{66}(t) < 1$.

Proposition 1 summarizes the above (intuitive) findings.

Proposition 1 *IPv6 adoption by Internet Content Providers (ICPs) is influenced as follows by connectivity quality*

- *ICPs have no incentives to become IPv6 accessible as long as translation gateways offer IPv6 users the same*

or better connectivity quality than native IPv6, i.e., $\beta(t) = 0$, *if* $q_{66}(t) \leq q_{64}(t)$;

- *Once native IPv6 connectivity quality exceeds that of translation gateways, the number of ICPs that choose to become IPv6 accessible grows with the number of IPv6 users,* i.e., $\frac{\partial \beta}{\partial x_6} > 0$ *if* $q_{66}(t) > q_{64}(t)$;

- [**model 1 only**] *As long as IPv6 connectivity is worse than IPv4 connectivity, allocating IPv6 addresses to current IPv4 users negatively impacts IPv6 adoption by ICPs,* i.e., $\frac{\partial \beta}{\partial \alpha} < 0$ *if* $q_{66}(t) < 1$;

- *When IPv6 connectivity quality is higher than that of IPv4, allocating IPv6 addresses to current IPv4 users always improves IPv6 adoption by ICPs,* i.e., $\frac{\partial \beta}{\partial \alpha} > 0$ *if* $q_{66}(t) > 1$.

3.2 Translation Traffic Growth

Section 3.1 focused primarily on the ICPs' IPv6 adoption decisions as captured by the parameter $\beta(t)$. In this section, we turn our attention to ISPs, and in particular the extent to which they can control the growth of translation traffic $T(t)$. This is of concern to ISPs because it affects the provisioning and, in turn, costs of translation gateways. In addition, one of the main motivations for ISPs to choose IPv6 is to eventually migrate to a translation-free Internet. Understanding the likelihood of such an outcome is, therefore, of interest.

The volume of translation traffic can be readily obtained from Eqs. (1) and (5). In this section, we investigate how different choices (by ISPs) of the decision variables $\alpha(t), q_{64}(t),$ and $q_{66}(t)$ affect $T(t)$, and in particular the ISP's ability to ensure $T(t) \leq a$, where a is a measure of the provisioned translation capacity. In other words, can the ISP "control" the volume of translation traffic through its three decision variables, $\alpha(t), q_{64}(t),$ and $q_{66}(t)$, knowing how they affect ICPs' decisions as predicted by Eq. (5). Note that this control is indirect, *i.e.,* through its influence on ICP decisions, and is not based on throttling or rate-limiting the amount of translation traffic that IPv6 users originate.

We consider three different cases for Eq. (5).

I. $q_{66}(t) < q_{64}(t) \leq 1.$

From Eq. (5) this results in $\beta(t) = 0$, *i.e.,* no ICPs make themselves IPv6 accessible. Translation traffic volume is, therefore, directly proportional to the number of new IPv6 users, $x_6(t)$, and the number of existing IPv4 users to which the ISP has also provided an IPv6 address, $\alpha(t)x_4$. Under model 1, reducing translation traffic is then best achieved by setting $\alpha(t) = 0$[7]. In summary, when the ISP is unable to provide IPv6 connectivity quality that is better than what is achievable through translation, it has little or no control on the volume of translation traffic, which stays below a only as long the number of IPv6 users itself remains below a, *i.e.,* $x_6(t) \leq a$. In other words, in **Configuration I**, as soon as the number of new IPv6 users exceeds a, the ISP has no choice but to upgrade its translation capacity to keep up with the growth of translation traffic. Alternatively, it could consider improving its support for IPv6, *e.g.,* by giving precedence to IPv6 traffic in its network and/or selecting more efficient routes, to improve $q_{66}(t)$ and entice more ICPs to become IPv6 accessible.

The next configuration considers the case where the ISP has improved IPv6 connectivity quality beyond that achiev-

able through translation, but not yet so that it surpasses IPv4 quality, *i.e.,* $q_{64}(t) \leq q_{66}(t) < q_{44} = 1$.

II. $q_{64}(t) \leq q_{66}(t) < 1.$

As in the previous scenario, the ISP's best strategy (under model 1) to minimize translation traffic volume is again to set $\alpha(t) = 0$.

There are, however, additional conditions that need to be satisfied to keep translation traffic volume below a, as the number of IPv6 users keeps growing. In particular, IPv6 connectivity quality, $q_{66}(t)$, needs to exceed translation quality, $q_{64}(t)$, by a certain margin. We explore this issue next. From Eq. (5) and assuming $q_{64}(t) \leq q_{66}(t) < 1$ (and $\alpha(t) = 0$), $q_{66}(t)$ and $q_{64}(t)$ must satisfy the following relation to ensure $T(t) \leq a$:

$$q_{66}(t) \geq q_{64}(t) + \frac{c_6}{x_6(t)} \left(1 - \frac{a}{x_6(t)} \right). \qquad (6)$$

In other words, if enough ICPs are to become IPv6 accessible to keep translation traffic below the provisioned capacity a, a minimum quality gap must exist between IPv4 and IPv6. For a given value of $q_{64}(t)$, the right-hand-side of Eq. (6) has a maximum for $x_6(t) = 2a$, which based on Eq. (1) is also where translation traffic peaks (at a value of a). Hence, to keep translation traffic below a for all values of $x_6(t)$, the minimum quality gap between IPv6 and translation gateways must satisfy:

$$\Delta_{64} = q_{66}(t) - q_{64}(t) \geq \frac{c_6}{4a}. \qquad (7)$$

As expected, Δ_{64} grows with c_6, the cost of IPv6 configuration for ICPs, and decreases with a, the provisioned capacity of translation devices. When combined with Eq. (5) and keeping $\alpha(t) = 0$, Eq. (7) indicates that maintaining such a quality gap is also sufficient to eventually ensure a complete migration to an IPv6 Internet (all ICPs have become IPv6 accessible). This occurs after the population of IPv6 users has grown large enough ($x_6(t) \geq 4a$) to ensure that the added cost of configuring IPv6 is justified for all ICPs.

Eq. (7) when combined with the conditions of **Configuration II** ($q_{64}(t) \leq q_{66}(t) < 1$) also has a more subtle implication. Specifically, it indicates that to keep translation traffic bounded without having to make IPv6 connectivity *better* than the current Internet, translation quality must be "bad enough," *i.e.,*

$$q_{64} < 1 - \frac{c_6}{4a}. \qquad (8)$$

When this is satisfied, the low translation quality is sufficient to entice more ICPs to become IPv6 accessible as the number of IPv6 users increases. How bad is bad enough depends on the ratio $\frac{c_6}{4a}$. When $\frac{c_6}{4a}$ is large, *e.g.,* because of high IPv6 configuration costs or limited translation capacity, this may translate into an unrealistic constraint, namely, an unacceptably poor (to users) translation capacity. Especially since one can expect IPv6 users to demand connection quality that is comparable to that of the IPv4 Internet, *i.e.,* $q_{64}(t) \approx q_{44} = 1$. When the constraint of Eq. (8) cannot be satisfied, the ISP has no choice but to make IPv6 connectivity equal or even better than that of the IPv4 Internet, if it is to keep translation traffic below the provisioned capacity a. This is the scenario we explore next, and we start by investigating when this may occur.

Combining the constraint $q_{66}(t) < 1$ (IPv6 connectivity remains worse than that of the IPv4 Internet) with Eq. (6)

---

[7] This is not an issue under model 2, where assigning IPv6 addresses to existing IPv4 users does not affect the volume of translation traffic.

yields the following inequality

$$f(x_6) = (1 - q_{64})x_6^2 - c_6 x_6 + ac_6 > 0 \,. \tag{9}$$

When $q_{64} \geq 1 - \frac{c_6}{4a}$, Eq. (9) has two real-valued roots of the form

$$x_6^{(1)} \quad = \quad \frac{c_6 - \sqrt{c_6^2 - 4ac_6(1 - q_{64})}}{2(1 - q_{64})} \tag{10}$$

$$x_6^{(2)} \quad = \quad \frac{c_6 + \sqrt{c_6^2 - 4ac_6(1 - q_{64})}}{2(1 - q_{64})} \tag{11}$$

As q_{64} varies in the range $(1 - \frac{c_6}{4a}, 1)$, the first root varies from $2a^-$ down to a^+, while the second root varies from $2a^+$ to ∞. Furthermore, $f(x_6)$ is negative in the range $(x_6^{(1)}, x_6^{(2)})$ and positive outside. This implies that Eq. (6) is satisfied only in the range $x_6(t) \in [a, x_6^{(1)})$. When $x_6(t)$ exceeds $x_6^{(1)}$, it becomes necessary for IPv6 quality to exceed that of IPv4, i.e., $q_{66}(t) \geq 1$, to ensure that the volume of translation traffic remains below a.

In other words, when translation quality is sufficiently good ($q_{64} \geq 1 - \frac{c_6}{4a}$), IPv6 adoption by ICPs is delayed to the point that IPv6 quality ultimately has to improve beyond that of IPv4 to keep translation traffic below a. This transition happens once the number of IPv6-only users, $x_6(t)$, reaches $x_6^{(1)}$, with the transition occurring earlier ($x_6^{(1)}$ decreases from $2a$ to a) as translation quality is better (approaches IPv4 quality). **Configuration III** explores this scenario in greater details.

III. $q_{66}(t) \geq 1$.

In this configuration, we are interested in how much better than the current Internet IPv6 connectivity should be to keep the volume of translation traffic below provisioning levels. From Eqs. (1) and (5), we obtain the following condition for $T(t) \leq a$

$$(q_{66} - q_{64})x_6^2 - (c_6 - \alpha x_4(q_{66} - 1))x_6 + ac_6 \geq 0 \,, \tag{12}$$

where for ease of notation we omitted dependency on t.

Eq. (12) implies

$$q_{66}(t) \geq \frac{q_{64}(t)x_6^2(t) + (c_6 + \alpha(t)x_4)x_6(t) - ac_6}{x_6^2(t) + \alpha(t)x_4 x_6(t)} \tag{13}$$

Note that when $\alpha(t) = 0$, Eq. (13) simplifies to

$$q_{66}(t) \geq q_{64}(t) + \frac{c_6}{4a} \tag{14}$$

This is consistent with Eq. (7).

As expected, higher values of c_6 (high configuration costs for ICPs) and lower values for a (limited translation capacity) call for a correspondingly better IPv6. Conversely, we easily see from Eq. (13) that increasing $\alpha(t)$ allows the inequality to be met with a smaller value of $q_{66}(t)$ for all $x_6(t)$. In other words, when IPv6 connectivity quality is better than that of IPv4, making more IPv4 users IPv6 capable is beneficial to keeping the volume of translation traffic low. This is because it makes for a larger number of users that can take advantage of the higher quality, which in turn increases the incentives for ICPs to become IPv6 accessible. The more ICPs are accessible over IPv6, the lower the volume of translation traffic.

More generally, from Eq. (13) we can compute the value

of $x_6^*(t)$ for which $q_{66}(t)$ realizes its maximum value, namely

$$x_6^*(t) = \frac{ac_6 + \sqrt{a^2 c_6^2 + ac_6 \alpha(t)x_4 \left[c_6 + \alpha(t)x_4(1 - q_{64}(t))\right]}}{c_6 + \alpha(t)x_4(1 - q_{64}(t))} \tag{15}$$

The specific expression for $x_6^*(t)$ of Eq. (15) does not add new insight, but it is worth noting that once $x_6(t)$ exceeds $x_6^*(t)$, it is possible for $q_{66}(t)$ to start decreasing again without risking an increase in translation traffic beyond a. As a matter of fact, it is even possible, although obviously neither necessarily practical nor desirable, to lower IPv6 quality back below that of IPv4 (but not below that of translation devices). This is because once the IPv6 user base is large enough, this alone is sufficient to entice enough ICPs to become IPv6 accessible.

The main findings from the above discussion are summarized in Proposition 2.

Proposition 2 *ISPs can control the volume of IPv6 traffic that undergoes translation (to IPv4) by adjusting the relative connectivity quality of IPv6 and translation compared to that of IPv4. Some of the trade-offs this involves are as follows*

- *As long as native IPv6 connectivity is worse than that of translation gateways, ICPs have no motivation to become IPv6 accessible. Hence, translation traffic keeps growing with the number of new IPv6 users.*
- *When translation quality can remain low, this alone is sufficient incentive for ICPs to become IPv6 accessible and keep translation traffic bounded, even with IPv6 quality below that of the IPv4 Internet.*
- *Keeping translation traffic bounded may require IPv6 connectivity to be better than that of the IPv4 Internet, when translation gateways are of high quality and/or limited capacity.*
- *When IPv6 quality exceeds that of the IPv4 Internet, allocating IPv6 addresses to current IPv4 users will hasten IPv6 adoption by ICPs and help keep the volume of translation traffic low.*
- *Once there are enough IPv6 users, this alone is sufficient incentive for ICPs to continue becoming IPv6 accessible and keep translation traffic bounded, even if IPv6 quality is no better than that of the IPv4 Internet.*

3.3   Numerical Examples

This section provides a few examples that illustrate the findings of the previous sections. It assumes $x_4 = 1$, *i.e.,* the number and traffic of IPv4 users is normalized to 1, $a = 0.1$ and $c_6 = 0.1$. In other words, translation devices have been provisioned to handle a traffic volume equal to 10% of the current IPv4 traffic, and the cost of IPv6 configuration is in the worst case ($\theta_i = 1$) equal to 10% of an ICP's revenue. Given these values, we consider two possible translation quality values, $q_{64} = 0.74$ and 0.79, where the first value satisfies Eq. (8), namely, $q_{64} = 0.74 < 1 - \frac{c_6}{4a} = 0.75$, but the latter doesn't. The impact of this difference is illustrated by comparing Figs. 1 and 2, with the former exhibiting consistently lower volumes of translation traffic. Additionally, Fig. 1 also shows that as indicated by Eq. (7), it is possible to keep $T(t) \leq a = 0.1$ with an IPv6 quality no better than that of IPv4 (Eq. (7) states that $\Delta_{64} \geq 0.25$, which calls for $q_{66} \geq 0.99$, and the figure plots translation traffic for $q_{66} = 1 > 0.99$). Figs. 1 and 2 also illustrate that improving
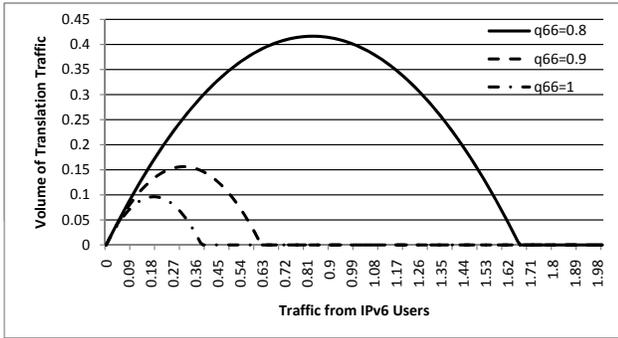
Figure 1: Impact of IPv6 Quality on Translation Traffic Volume ($q_{64} = 0.74$).

translation quality calls for improving IPv6 quality if one is to keep translation traffic below the provisioned capacity. A similar conclusion applies to ensuring full migration of the IPv4 Internet to IPv6, *e.g.,* when $q_{64} = 0.74$, β reaches 1 for $x_6 = 1.67$ and $q_{66} = 0.8$, but increasing q_{64} to 0.79 calls for correspondingly increasing q_{66} to 0.85 to achieve the same result.
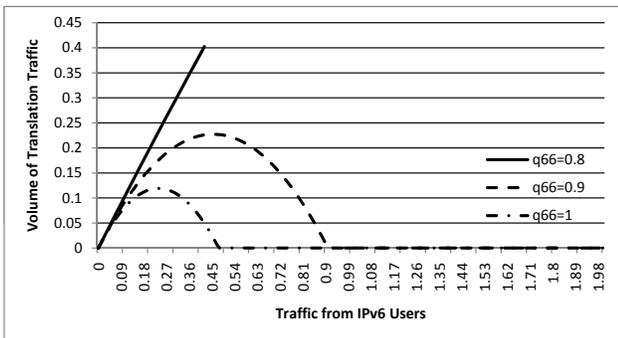

Figure 2: Impact of IPv6 Quality on Translation Traffic Volume ($q_{64} = 0.79$).

Another perspective on the impact of q_{64} on q_{66} is illustrated in Fig. 3 that plots as a function of the number of IPv6 users, the minimum required IPv6 quality to keep $T(t) \leq a = 0.1$. The figure also shows that once the number of IPv6 users is large enough, *i.e.,* exceeds the value of Eq. (15), translation traffic remains bounded even when IPv6 quality is no better than that of IPv4. As the figure indicates, it is even possible, although as mentioned earlier neither desirable nor practical, to decrease IPv6 quality below that of IPv4 without risking exceeding the capacity of translation devices.

4. CONCLUSION

This paper used a simple model to explore how quality and capacity of translation devices and IPv6 quality could affect both migration of the IPv4 Internet to IPv6, and the volume of traffic that translation devices need to handle.

In spite of its simplicity and obvious limitations, the model helped elucidate a number of interesting issues, and in particular the ambiguous role of translation gateways. Those devices are mandatory to let IPv6 users access the current
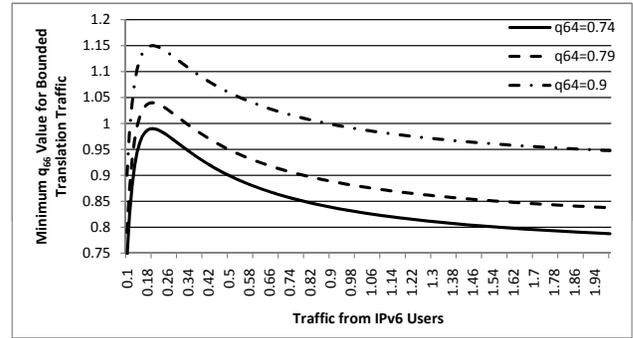

Figure 3: Minimum IPv6 Quality to Keep Translation Traffic Volume Below Provisioned Capacity.

Internet that is reachable mostly only over IPv4. The quality of those devices must, therefore, be high enough to satisfy those users. On the other hand, if their quality is too high, it will not only slow down an eventual migration to IPv6 (because ICPs will have less of an incentive to become IPv6 accessible), it may also require that IPv6 quality exceeds that of the IPv4 Internet if translation traffic is to be kept bounded as the number of IPv6 users grows. Another more intuitive finding is the role of IPv6 connectivity quality. Until native IPv6 connectivity is of a higher quality than what translation gateways offer, ICPs have no incentives to become IPv6 accessible, and the volume of translation traffic will keep growing. Conversely, high-quality IPv6 connectivity can alone be a sufficient incentive for ICPs to become IPv6 accessible early on, and therefore help keep the volume of translation traffic low.

These findings can be translated into two simple recommendations to ISPs that have adopted IPv6:

1. Do not make translation gateways any better than you have to;
2. Make IPv6 connectivity quality as high as possible and preferably higher than that of the current IPv4 Internet; at least in the early phases when the number of IPv6 users is low.

The first recommendation ensures that the incentives for ICPs to avoid becoming IPv6 accessible and force IPv6 users to access them through translation gateways are as low as possible. The second recommendation has the same goal but seeks to realize it in opposite ways, namely, encourage ICPs to become IPv6 accessible as early as possible because of the better quality of IPv6 rather than the low quality of translation devices. This is especially important in the early stages, when the small number of IPv6 users by itself offers ICPs little incentive to become IPv6 accessible.

There are obvious caveats to the above recommendations. The least of which is the distributed nature of the Internet that is made-up of a collection of interconnected autonomous entities. The failure of a few of them to offer high-quality IPv6 connectivity could affect end-to-end IPv6 quality for many users and ICPs. Nevertheless, if the larger ISPs, which to some extent stand to gain the most from a faster migration to IPv6, lead the way in offering high(er) quality IPv6 connectivity, this may be sufficient incentive (the corresponding user base is large) for many ICPs to consider making themselves IPv6 accessible.

There are many extensions one could consider to make the simple model of this paper more realistic. For example, one could incorporate aspects of competition among ICPs and possibly ISPs, and explore how this might affect IPv6 adoption decisions. Other extensions of interest include allowing heterogeneity in content popularity and revenues as well as in connectivity quality, and minimizing the model's reliance on exogenous parameters. One such possible direction is to include users among the decision makers, *i.e.,* endogenize $x_6(t)$ by making it a function of the connectivity quality that new users experience. Additionally, it may be of interest to capture more accurately the tension between translation costs and connectivity costs, *i.e.,* formulate a joint optimization that will incorporate both cost components and seek to characterize the strategy that yields the lowest overall cost.

Another important extension, is to offer empirical insights into how migration to IPv6 is taking place. Real-world data can help quantify both IPv6 adoption patterns as well as differences in connectivity quality, and can be used to validate and extend the results from our model on the interaction between connectivity quality and adoption. Conversely, investigating the reasons behind those differences in quality can help pinpoint problem areas and what ISPs need to focus on if they are to follow recommendation 2. These are the motivating factors behind the previously mentioned measurement efforts whose partial results are available at http://mnlab-ipv6.seas.upenn.edu/monitor/index.html. We hope to be able to report on those more completely in the near future.

Acknowledgments

5. REFERENCES

[1] Alexa Top Web Sites. Available at http://www.alexa.com/topsites.

[2] S. Cheshire. IPv6 Adoption - Keynote presentation at IETF 72 Technical Plenary. Available at http://www.stuartcheshire.org/IETF72/, July 2008.

[3] K. Das. US Government using IPv6. Available at http://ipv6.com/articles/general/ US_Government_IPv6.htm.

[4] S. Deering and R. Hinden. Internet Protocol, version 6 (IPv6) specification. RFC 1883, December 1995.

[5] R. Draves. Default address selection for Internet Protocol version 6 (IPv6). RFC 3484, February 2003.

[6] A. Durand, Ed. Dual-stack lite broadband deployments post IPv4 exhaustion. Internet Draft (draft-ietf-softwire-dual-stack-lite-02), October 2009. (Work in progress).

[7] B. Edelman. Running out of numbers: Scarcity of IP addresses and what to do about it. In *Proc. First Conference on Auctions, Market Mechanisms and Their Applications*. Springer-Verlag's Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, 2009.

[8] G. Huston. Is the transition to IPv6 a "market failure"? Available at http://www.potaroo.net/ispcol/2009-09/v6trans.html, September 2009.

[9] C. Labovitz. The end is near, but is IPv6? Available at http://asert.arbornetworks.com/ 2008/08/the-end-is-near-but-is-ipv6/, August 2008.

[10] Technical and economic assessment of Internet Protocol version 6 (IPv6). US Department of Commerce - Available at http://www.ntia.doc.gov/ntiahome/ntiageneral/ ipv6/final/ipv6final.pdf, January 2006.

---

[8]See http://www.caida.org/workshops/wie/0909/.