Modeling the Adoption of new Network Architectures

Dilip Joseph Nikhil Shetty John Chuang Ion Stoica
{dilip@cs, nikhils@eecs, chuang@ischool, istoica@cs}.berkeley.edu
University of California at Berkeley

ABSTRACT

We propose an economic model based on user utility to study the adoption of new network architectures such as IPv6. We use mathematical analysis and simulation studies to understand the role of various factors such as user and network benefits, switching costs, and the impact of converters on the adoption of new network architectures. In addition to corroborating various commonly held beliefs about new network architecture adoption, our analysis and simulation studies also reveal several surprising and non-intuitive results. For example, while in general, increasing the efficiency of converters hastens the adoption of new network architectures, there are cases in which more efficient converters hinder the adoption of such architectures. This and other results in the paper increase our understanding of new network architecture adoption and guide the design and implementation of mechanisms to hasten new network architecture adoption.

1. INTRODUCTION

Despite being developed over a decade ago, IPv6 still has not achieved widespread adoption. Neither the looming threat of IPv4 address exhaustion, IPv6's close resemblance to IPv4 nor the widespread availability of dual stack IPv4+IPv6 operating systems has spurred IPv6 onto the mainstream Internet [2]. Meanwhile, the networking research community has proposed many more network architectures [14, 7, 15, 10, 9, 13, 3], some radically different from IPv4 and some without clear transition and deployment mechanisms. What factors will aid widespread adoption of these new network architectures? In this paper, we propose an economic model to study the adoption of new network architectures.

Our model is based on the *user utility* concept. A user represents a single individual or an entire organization. A user of a particular network architecture receives standalone benefits which are unaffected by the presence or absence of other users, as well as network benefits arising from the ability to communicate with other users of the same architecture. A user can switch to a new network architecture that offers better utility or can adopt a *converter* that provides partial benefits of the new architecture while remaining with the old architecture. We mathematically analyze the model from the standpoint of aggregate utility of all users in the system in order to understand the impacts of new network architecture adoption on the society as a whole. As the decision to adopt a new technology is in practice made by an individual user, we also study the system dynamics from the perspective of an individual user via mathematical analysis and simulations.

Our analysis and simulation results reveal both obvious, as well as surprising insights about the adoption of new network architectures. We describe the following and other results in detail in Sections 4 and 5:

- New network architecture adoption needs to withstand a period of decreasing total system utility till a critical mass of users is reached. Incentives from government or industry can encourage users to adopt the new architecture and expedite attainment of critical mass.
- Converters aid new network architecture adoption by reducing the loss in total utility to be endured till critical mass is attained. However, converters may be detrimental to complete adoption unless they are carefully designed and engineered.
- Adoption of a new network architecture happens faster if users get the news about other users switching more quickly. However, when highly efficient two-way converters are available, quick dissemination of news about the progress of new network architecture adoption may hinder adoption.
- New network architecture adoption stalls if network effects do not fall within an upper and lower bound determined by current network conditions. High network effects tightly lock users to the current network architecture, while low network effects offer no incentive to incur switching costs and adopt a new architecture.

*CoNEXT'*07, December 10-13, 2007, New York, NY, U.S.A.

This work is only an initial step in studying the adoption of new network architectures from an economic standpoint. Our current economic model is very basic. The parameter values used in the analytical and simulation studies do not directly map on to real world numbers. We discuss these and other limitations of our work in Section 7.

2. RELATED WORK

Adoption of new technologies and products has been extensively studied in Economics [8, 6]. Adoption of new network architectures is similar to the adoption of any new technology in many ways - for example, switching costs and network benefits are important to both. However, there are some important differences. In most new technology adoption scenarios, there are multiple organizations competing with each other to further one particular technology. The adoption of the new technology depends on how these organizations compete with each other on price and features. For new network architectures, especially IPv6, there are currently no opposing organizations pushing different technologies. Opposition to a new network architecture will come from organizations unwilling to foot the switching costs.

There have been few papers which study the adoption of new network architectures. [2] estimates the progress and costs of IPv6 adoption based on interviews with infrastructure providers, application vendors, ISPs and users. We propose a general user-focused model for new network architecture adoption and study it using mathematical analysis and simulations. The Internet Standards Adoption (ISA) framework proposed by [11] identifies *usefulness of features* and *environmental conduciveness* as the factors influencing the mode of adoption of a new Internet standard. These factors are similar to the standalone utility and network benefits considered in our model. However, [11] uses case studies to construct and illustrate the model and does not perform an analytical or simulation study. Unlike [2] and [11], our paper focuses on the role of converters in the adoption of new network architectures. In [5], the authors model and simulate the adoption of secure BGP protocols and define the switching threshold as an adoptibility metric. Our paper models adoption of generic new network architectures instead of a single class of protocols.

3. MODEL

Our model to study the adoption of new network architectures is based on the *utility* or benefits offered by the network architecture to a *user*. We believe that individual consumers and organizations, and killer applications enabled by new network architectures will be the key drivers for new network architecture adoption [2]. Hence, a *user* in our model represents an individual consumer or an organization, and not ISPs or infrastructure vendors. A user of a particular network architecture receives two types of benefits:

Standalone benefits which do not depend on the presence or absence of other users of the same architecture. For example, an IPv6 user derives standalone utility from the vast address space and automatic address configuration provided by IPv6.

Network benefits derived from the ability to communicate with other users of the same architecture. For example, *i3* [14] users benefit from the ability to communicate with all other *i3* users.

Our model consists of N users, each of whom has adopted either network architecture A or B. Network architecture A represents the incumbent architecture (for example, IPv4) and B represents the new architecture (for example, IPv6 or *i3*). Fraction x_A of the N users in our model are users of architecture A, while the fraction $x_B = 1 - x_A$ are B users. Table 1 describes the notation used in this paper.

An A user switches to B if B offers higher utility than A even after accounting for switching costs. Rather than making a complete switch from A to B, an A user may also choose to remain with A and use an AB converter. An AB converter provides a portion of the standalone and network utility offered by B to an A user. An IPv4-IPv6 gateway and client side software like the Hexago Gateway6 client [1] are examples for IPv4-IPv6 converters. OCALA [12] is a generic converter for new network architectures. A fraction x_{AB} of A users run AB converters; A fraction x_{BA} of B users run BA converters. We assume that converters are two-way, i.e., an AB converter enables an A user to communicate with all B users and also enables all B users to communicate with an A user running an AB converter.

The utility enjoyed by an A user who does not use an AB converter, i.e. an AONLY user, is given by:

$$U_{Aonly} = \alpha_A + \beta N x_A + \beta N x_{BA} x_B (1 - q_A) \qquad (1)$$

α_A is the standalone benefit provided by A. $\beta N x_A$ is the network benefit due to the ability to communicate with the $N x_A$ A users. For model simplicity and ease of analysis, we have assumed the commonly used linear model of network benefits [8] enjoyed by an individual user[1], with a single parameter β controlling the importance of the network effects. Our ongoing work considers alternate models of network benefits like the logarithmic model [4]. $\beta N x_{BA} x_B (1 - q_A)$ is the network effect benefit due to the $N x_{BA} x_B$ B users who have adopted BA converters. A BA converter does not offer full compatibility with A. Hence an A user communicating with a B user who has adopted a BA converter receives only a fraction $(1 - q_A)$ of the network benefits of communicating with an A user. The parameter q_A measures the efficiency of a BA converter.

The utility enjoyed by an A user who uses an AB converter, i.e. an AB user, is given by:

$$U_{AB} = (1 - r_A)(\alpha_A + \beta N x_A) + \beta N x_B (1 - q_B) + t_B \alpha_B \qquad (2)$$

The first term represents the benefits an AB user gains from network architecture A. The sum of the two remain-

---

[1] This agrees with Metcalfe's Law: *The value of a telecommunication system is proportional to the square of the number of users in the system.*

A	Incumbent network architecture (e.g. IPv4)
B	New network architecture (e.g. IPv6)
A(B) user	A user of architecture A(B) with or without an AB(BA) converter
AONLY(BONLY) user	An A(B) user not using an AB(BA) converter
AB(BA) user	An A(B) user who uses an AB(BA) converter
N	Total number of users in the system
β	Parameter controlling the magnitude of network effects
x_A (x_B)	Fraction of the N users who are A (B) users
x_{AB} (x_{BA})	Fraction of the A (B) users with AB (BA) converters
α_A (α_B)	Standalone utility offered by architecture A (B)
$(1-q_A)$ $((1-q_B))$	Fraction of network benefits of A (B) offered by a BA (AB) converter
r_A (r_B)	Degradation in utility offered by A(B) due to AB (BA) converter use
t_A (t_B)	Fraction of A(B) standalone utility provided by a BA (AB) converter
S	Generic term used for switching cost between AONLY & AB, AONLY & BONLY, BONLY & BA

Table 1: Notation used in the paper.

ing terms represents the benefits obtained from network architecture B, via the AB converter. $(\alpha_A + \beta N x_A)$ is the utility offered by network architecture A to an AB user, irrespective of B users. This utility is sometimes degraded by the use of an AB converter. For example, running an IPv4-IPv6 converter may expose the user to higher security risks due to outdated firewall rules. A user running the OCALA proxy to communicate with *i3* users may experience slightly increased latencies for regular IPv4 communication due to the packet interception and processing performed by the OCALA proxy. The parameter r_A captures the potential degradation caused by an AB converter on the utility offered by A.

In equation 1, we saw that an AONLY user receives partial benefits of communicating with B users who have adopted BA converters. Through an AB converter, an AB user is able to communicate with all B users (both BONLY and BA). Since converters are not 100% efficient, as mentioned earlier, the network benefit obtained by an AB user is degraded by a factor $(1 - q_B)$, where q_B captures AB converter efficiency. For simplicity, we treat all B users (both BONLY and BA) alike and apply a single degradation $(1 - q_B)$ on the network benefit due to B users contactable via the AB converter.

In addition to the ability to communicate with B users, an AB converter can also provide an A user with a portion of the standalone benefits offered by B. For example, the OCALA proxy provides *i3*-style mobility support to a user's applications, in addition to enabling communication with other *i3* users. The term $t_B\alpha_B$ in Equation 2 represents this benefit, where t_B is the fraction of the standalone benefits of B provided by an AB converter.

The expressions for utility received by BONLY and BA users are similar to equations 1 and 2 respectively. The total utility received by all the N users in our model is given by:

$$\begin{aligned} TU = N(1-x_{AB})x_A U_{Aonly} + N x_{AB} x_A U_{AB} + \\ N(1-x_{BA})x_B U_{Bonly} + N x_{BA} x_B U_{BA} \end{aligned} \tag{3}$$

Our economic model for new network architecture adoption is very basic. We have made the simplifying assumptions described above in order to keep the model simple and tractable for mathematical analysis. The generality of our model may prevent it from capturing the impact of features that are very specific to a certain new network architecture. The limitations of our model are described in detail in Section 7. Although our model is very basic, it enhances our understanding of new network architecture adoption by corroborating some of the commonly held beliefs and by shedding light on some non-obvious aspects of new network architecture adoption, through the mathematical analysis and simulations described in the next two sections.

4. MATHEMATICAL ANALYSIS

In this section, we analyze the model formulated in the previous section and quantify the impact of various parameters on the adoption of the new network architecture B. We first study new network architecture adoption based on the total utility of all users in the system. We then analyze the model from the perspective of individual users making the switching decision purely to maximize individual utility.

4.1 Total Utility

Studying the model from a total utility standpoint helps us understand the impacts of new network architecture adoption on the society as a whole. Maximizing the aggregate utility enjoyed by all users is a desirable social goal[2]. Maximizing aggregate utility usually requires coordinated action by all users. One way to attain coordinated action is through government mandate. The total utility of the system depends on the number of users of each kind. We simplify our analysis into 3 distinct cases: (1) AONLY and BONLY users, (2) AB and BONLY users, and (3) BA and AONLY users.

---

[2] Except in cases where there is gross inequality in distribution of utility among different users.

Our analysis of new architecture adoption from a total system utility standpoint leads to three main observations:

1. There is a period of decreasing total system utility associated with the adoption of a new network architecture. This result is true even when the new architecture being adopted has higher standalone utility than the current architecture. In some cases, converters help reduce or eliminate this period of decreasing utility. A critical penetration level of new network architecture users is needed to get adoption rolling and this may hence require government intervention or economic incentives.
2. Both AB and BA converters aid the adoption of the new network architecture B. Increasing the functionality and efficiency of BA converters hastens adoption. However, complete adoption of B is hindered if AB converters are "too good". For organizations wishing to enable the move to the next architecture, this advocates a controversial strategy of intentionally keeping the standalone utility offered by AB converters low in order to promote complete adoption of the new network architecture B.
3. New network architecture adoption becomes harder when the importance of network effects or the number of users in the system increases.

In the remainder of this section, we describe these and other results in detail.

4.1.1 *Case 1:* AONLY *and* BONLY *users*

When no AB or BA converters are available, all users in the system are of type AONLY or BONLY (i.e. $x_{AB} = x_{BA} = 0$). In this case, the total utility of the system hits a minimum point, or a trough, when the fraction of B users equals $x_B^* = \frac{1}{2} + \frac{\alpha_A - \alpha_B}{4N\beta}$. Figure 1 plots the total utility curves for different values of $\frac{\alpha_B}{\alpha_A}$ and different penetration levels of B. When A and B offer identical standalone utilities, the total utility of the system is minimized when there are equal number of A and B users.

In the initial B adoption phase (when x_B is low), adoption is infeasible from a total utility standpoint. For example, when $\alpha_A = \alpha_B$, the system as a whole has to bear a utility hit of up to 25% of the current utility in order to go past 50% B penetration. If the system can be coerced to go beyond 50% B penetration, adoption of B becomes feasible and proceeds automatically as the total utility keeps increasing when users switch from A to B. Hence, overcoming the critical penetration threshold is crucial. Government intervention to coerce A users to switch to B and economic incentives to offset the initial loss in utility can aid in reaching the critical penetration level. The critical penetration level required and the loss in utility decrease as the relative superiority of B over A increases. Figure 1 shows that the penetration threshold is 37.5% and the maximum utility loss is 14% when the standalone utility of B is 1.5 times that of A. Thus, enhancing a new network architecture with features that increase its standalone utility relative to the incumbent architecture will hasten its adoption.

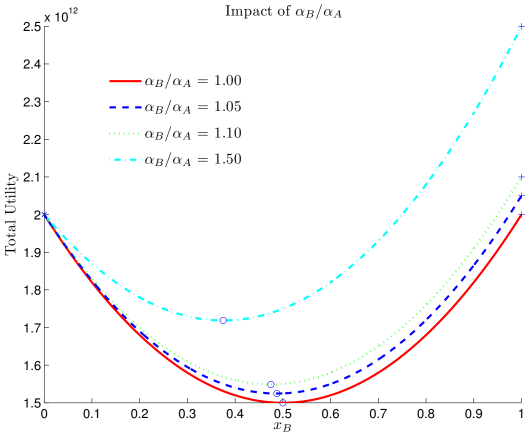

Figure 1: Total utility for different values of $\frac{\alpha_B}{\alpha_A}$

Network effects also play a key role in hastening or slowing down the adoption of a new network architecture. Figure 2 plots the total utility curves for different levels of network effects (captured by the model parameters $N\beta$) and different penetration levels of B, when $\frac{\alpha_B}{\alpha_A}$ = 1.5. The graphs in Figure 2 illustrate the expected result that the total utility increases as the importance of network effects increases. Figure 2 also shows that adoption of B becomes harder as the importance of network effects increases, due to the following two reasons: (i) Higher network effects increase the critical mass of B users required to overcome the point of lowest total system utility and make adoption of B feasible. (ii) Higher network effects also increase the utility loss to be overcome while moving from A to B, both in terms of absolute numbers and as a fraction of the total utility at zero B penetration. Our total utility analysis here thus shows that adoption of B stalls if the network effects are very high. Later, in our simulation results that consider the perspectives of individual users, we show that adoption of B also stalls if the network effects are too low.

If we separately analyze the two parameters contributing the network effects, viz. N and β, Figure 2 shows that, for the same value of $N\beta$, N has a greater effect on the total utility than β. In terms of absolute numbers, the utility loss to be overcome for B adoption is higher in the case where N is higher. The increased total utility attained at complete B adoption is also higher. However, if measured as a percentage of the utility at zero B penetration, the utility loss as well as the increased utility at complete B adoption are identical for the same value of $N\beta$, irrespective of the individual values of N and β. This result is a direct consequence of the linear model for network effects considered in Section 3.

4.1.2 *Case 2:* AB *and* BONLY *users*

When only AB converters are available and $U_{AB} \geq U_{Aonly}$, all users in the system are of types BONLY or AB (ie. $x_{AB} =$

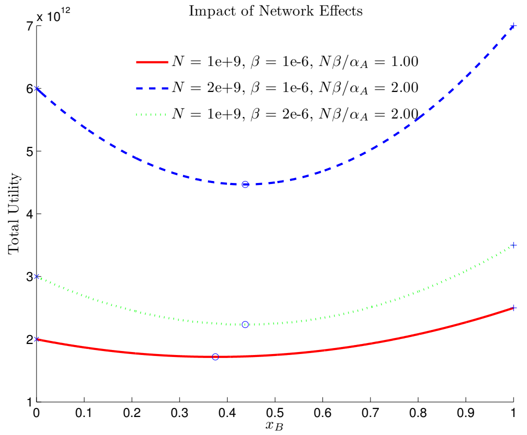

Figure 2: Impact of Network Effects ($\frac{\alpha_B}{\alpha_A}$ = 1.5)

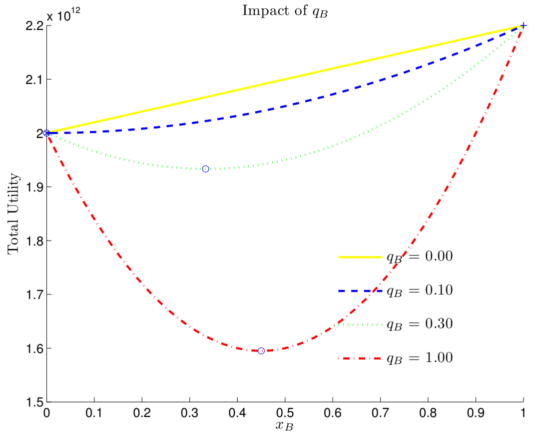

Figure 3: Impact of AB Converter Efficiency ($N = 10^9, \beta = 10^{-6}$, $\frac{\alpha_B}{\alpha_A}$=1.2)

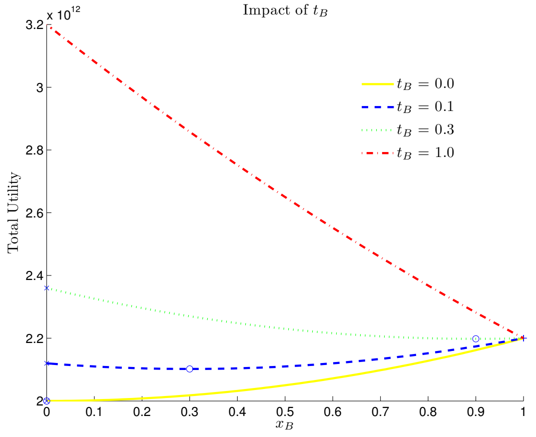

Figure 4: Impact of t_B ($N = 10^9$, $\beta = 10^{-6}$, $\frac{\alpha_B}{\alpha_A}$=1.2, q_B = 0.1, $r_A = 0$

1 and $x_{BA} = 0$). $U_{AB} \geq U_{Aonly}$ holds only if $x_B \geq \frac{r_A(\alpha_A+\beta N)-t_B\alpha_B}{\beta N(1-q_B+r_A)}$. This condition is trivially satisfied if the AB converter causes no degradation in the utility offered by A (i.e. r_A = 0). When $r_A < 2q_B$, the total system utility is minimized when

$$x_B^* = \frac{q_B - r_A}{2q_B - r_A} + \frac{\alpha_A(1-r_A) - \alpha_B(1-t_B)}{2\beta N(2q_B - r_A)} \qquad (4)$$

Let us assume that B offers 1.2 times the standalone benefits of A and that the AB converter does not degrade the standalone or network benefits associated with A and nor does it provide an A user with any of the standalone benefits of B (i.e., $\frac{\alpha_B}{\alpha_A}$ = 1.2, $r_A = 0$, $t_B = 0$). Figure 3 shows that the initial B penetration required to set the adoption of B automatically rolling decreases when the AB converter becomes more efficient (i.e. $q_B \rightarrow 0$). The loss in total utility to be overcome in moving from the low initial B penetration to complete B adoption also decreases with increasing AB converter efficiency. In fact, as seen in Figure 3, B penetration can be achieved without experiencing any decrease in total utility if the AB converter efficiency is greater than 90% (i.e. $q_B < 0.01$). Our analysis thus shows that, from a total utility standpoint, more efficient AB converters help in complete adoption of B. However, from an individual user standpoint, as described later in our simulation results, highly efficient AB converters are detrimental to the adoption of B.

If an AB converter provides an A user with a portion of the standalone benefits of B, i.e. $t_B > 0$, the initial B penetration required for making complete B adoption feasible increases. Figure 4 shows that for high values of t_B, complete B adoption is infeasible as the total utility at 100% B penetration is smaller than that at lower B penetration levels. There is no incentive for an A user to switch to B if an AB converter provides a sufficiently large portion of the standalone benefits of B without degrading performance.

A high value of t_B does not help adoption of B even when the AB converter imposes a heavy degradation on the standalone and network benefits associated with A ($r_A > 2q_B$), as shown in Figure 5. A high value of t_B simply increases the total utility at $x_B = 0$ and hence reduces the additional utility to be gained by moving to $x_B = 1$. This observation advocates a controversial strategy of intentionally keeping the B benefits offered by an AB converter low in order to promote complete adoption of B. For example, an AB converter may be programmed to occasionally drop packets destined to a B user. However, this may not be implementable in a market scenario free from government intervention and with competition to provide the best converter.

4.1.3 Case 3: BA *and* AONLY *users*

When only BA converters are available and $U_{BA} \geq U_{Bonly}$, all users in the system are of types AONLY or BA (i.e. $x_{AB} = 0$ and $x_{BA} = 1$). $U_{BA} \geq U_{Bonly}$ holds only when $x_B \leq \frac{\alpha_A t_A + N\beta(1-q_A)}{\beta N(1+r_B-q_A)}$. This condition trivially holds if the BA converter does not degrade the utility of B (i.e. $r_B = 0$). When $r_B < 2q_A$, the total system utility is minimized when

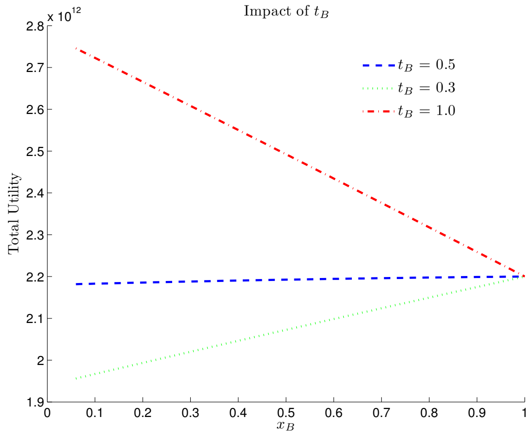

Figure 5: Impact of t_B ($N = 10^9, \beta = 10^{-6}, \alpha_B/\alpha_A = 1.2, q_B = 0.1, r_A = 0.21$). We consider only the range $0.06 \geq x_B \leq 1$ in order to ensure that $U_{AB} \geq U_A$ holds.

$$x_B^* = \frac{q_A}{2q_A - r_B} + \frac{\alpha_A(1 - t_A) - \alpha_B(1 - r_B)}{2\beta N(2q_A - r_B)} \tag{5}$$

The similarity of equations 5 and 4 implies that the impact of BA and AB converter efficiencies on the adoption of B follow similar trends: Higher the BA converter efficiency, easier is the complete adoption of B. In Case 2, we made a forward reference to simulation results which show that highly efficient AB converters impede the adoption of B. Later simulation results will highlight a similar observation about BA converters – increasing BA converter efficiency does not always aid the adoption of B. However, the cause behind this result is very different from that behind the result associated with AB converters.

In Case 2, we saw that high t_B values hamper the adoption of B. Contrastingly, in Case 3, we find that high t_A promotes the adoption of B (Figure 6). Increasing t_A raises the additional total utility attained as more and more users adopt B. Thus, in order to hasten the adoption of B, we should build BA converters that offer a substantial portion of the standalone benefits of A. We also need to take care to minimize any degradation caused by the converter on the standalone benefits of B. Figure 7 shows that when r_B is high (for example 10%), the total utility at 90% B penetration is lower than the total utility at lower levels of B penetration, and hence adoption of B is infeasible from a total utility standpoint.

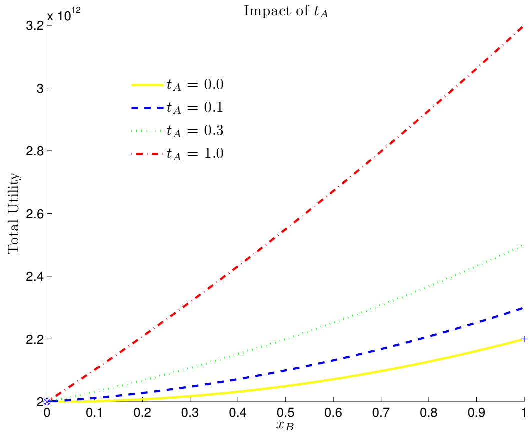

Figure 6: Impact of t_A ($N = 10^9, \beta = 10^{-6}$, $\frac{\alpha_B}{\alpha_A}= 1.2$, $q_A = 0.1, r_B = 0$

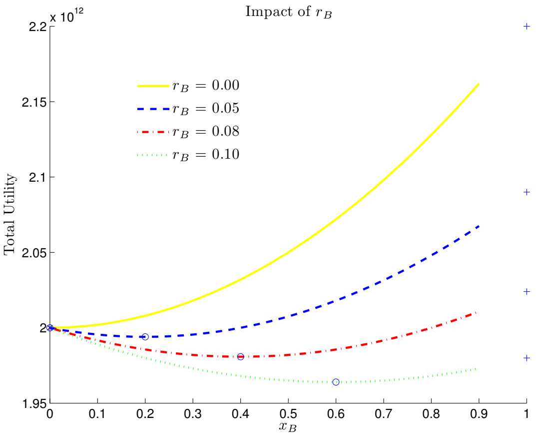

Figure 7: Impact of r_B ($N = 10^9$, $\beta = 10^{-6}$, $\frac{\alpha_B}{\alpha_A}$=1.2, $q_A = 0.1$). We consider only the range $x_B \leq 0.9$ in order to ensure that $U_{BA} \geq U_B$ holds.

4.1.4 *Trade-offs in Converter Design*

The analysis presented in this section can be used to study trade-offs associated with converter design. For example, should an IPv4-IPv6 converter be made more efficient at the expense of increasing degradation on IPv4 utility? Graphs like Figure 8 can be constructed to analyze situations like these if parameter values can be reasonably estimated, a currently difficult task due to lack of data. Figure 8 shows that a 90% efficient BA converter that causes 5% degradation offers higher total system utility with increasing B adoption than a 95% efficient BA converter that causes 10% degradation. Thus in this scenario, reducing converter degradation is more important for promoting B adoption than increasing converter efficiency.

4.2 Individual Utility

Most often, it is the end users that decide whether a given network architecture should be adopted or not. This decision is made in a selfish way by the users, without any consideration for the social utility of their actions. In this section, we will study the adoption of the new network architecture B from the standpoint of individual users. For simplicity, we will limit our analytical study to the three cases introduced in the previous section. Also, we will continue to use the existing model wherein all users are homogeneous, i.e., they have the same parameters in their utility functions.

The results in this section show that it is possible to make

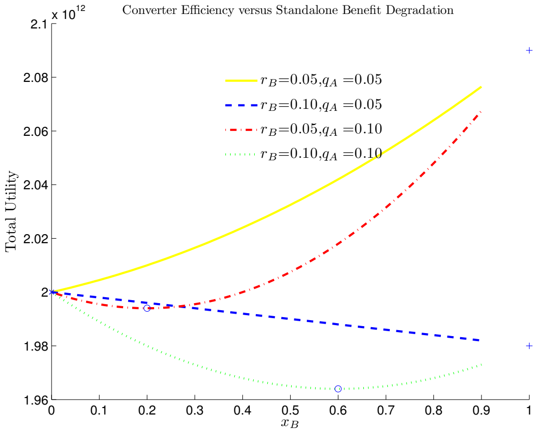

Figure 8: Trade-off between converter efficiency and degradation($N = 10^9, \beta = 10^{-6}$, $\frac{\alpha_B}{\alpha_A} = 1.2$). We consider only the range $x_B \leq 0.9$ in order to ensure that $U_{BA} \geq U_B$ holds.

the whole set of users switch to a new architecture by incentivizing a small portion of users. This fraction of users is smaller than the fraction of users required to take the total social utility past its minimum point. We also show that, as converter efficiency increases, the gap between the two above mentioned fractions widens.

4.2.1 *Case 1:* AONLY *and* BONLY *users*

If no converters are available, an A user switches to B if $U_{Bonly} \geq U_{Aonly} + S$, where S is the switching cost. When the B penetration level is $x_B = \frac{1}{2} + \frac{S-(\alpha_B-\alpha_A)}{2N\beta}$, the user is ambivalent between A and B. We call this value of x_B the *Equivalence Point*, x_B^e. A higher value of α_B lowers the equivalence point, i.e., a lower B penetration level is required to make switching to B attractive for an A user. This means that if the B penetration level rises beyond this level, then every user is automatically incentivized to switch over to the network architecture B. Once the penetration passes this fraction, all the users begin opting for architecture B *en masse*. This point can hence also be called a *tipping point*.

If we assume a zero switching cost, we find that $x_B^* = x_B^e + \frac{\alpha_B-\alpha_A}{4N\beta}$ (Section 4.1.1). Assuming $\alpha_B \geq \alpha_A$, this implies that we need a lower B penetration to entice A users to switch to B individually rather than to switch collectively. In other words, the number of B users needed to make all users shift to architecture B is lower than that required to exceed the minimum point in the total utility. This is interesting because it implies that we do not need to coerce as large a number of users to adopt the new architecture as we determined in the total utility case. Instead, by just incentivizing a smaller fraction of users to adopt the new architecture, we can cause the complete shift to the new architecture.

Intuitively, x_B^e is smaller than x_B^* because x_B^* takes into account the network effects between all pairs of users, while x_B^e is very myopic in scope. In other words, when B penetration is greater than x_B^e, even though each user is incentivized to shift over to B for his own selfish good, the social utility reduces. This is because a bigger fraction of users are presently using network architecture A. Any users shifting selfishly to architecture B causes a loss of utility to a large fraction of users (those using A), while increasing the utility of a lesser fraction (those using B), thus on a whole, reducing the total social utility. This effect reduces when the number of users becomes very large as the impact of a single user's shift from A to B becomes negligible. This can be observed by setting N to ∞ in the relationship between x_B^* and x_B^e. We also observe that the gap between x_B^* and x_B^e widens as the relative superiority of B over A increases.

4.2.2 *Case 2:* AB *and* BONLY *users*

In Case 1 above, we observed that if the penetration of B exceeds a certain level, x_B^e, then all users voluntarily opt for the new network architecture B. The presence of converters complicates the analysis in Case 2, by introducing a large number of model parameters and adoption choices. If we keep the fraction of B adopters fixed (which can happen due to slow adoption of B), we find that AONLY users are incentivized to adopt an AB converter whenever $U_{AB} \geq U_{Aonly} + S$, where S represents the costs of installation, learning and management of the converter. On substituting the values for U_{AB} and U_{Aonly} from Equations 1 and 2 and simplifying, we get the condition $x_B \geq \frac{S-\alpha_B t_B+\alpha_A r_A}{(1-q_B+r_A)\beta N} + \frac{r_A}{1-q_B+r_A}$. This threshold value for x_B, called x_{AB}^e, is the equivalence point between AONLY and AB. As N becomes very large, the first term in x_{AB}^e becomes negligible while the second term goes to $\frac{r_A}{1-q_B+r_A}$. If AB converters are very efficient (i.e. q_B is close to 0) and cause very little self-degradation (i.e., r_A is close to 0), then x_{AB}^e is close to 0. On the other hand, as seen earlier, the equivalence point between AONLY and BONLY is $x_B^e = \frac{1}{2}$ (assuming large N). x_B^e is in general greater than x_{AB}^e, since x_{AB}^e is close to zero, as argued earlier. When x_B is between x_{AB}^e and x_B^e, all A users are incentivized to take up the converter but not to adopt B. Hence, if B adoption is slow and is in the range (x_{AB}^e, x_B^e), all A users will adopt AB converters and we will have a situation where there are only AB and B users.

Let us make this idea concrete with the following example. Suppose $r_A = 0.1$ and $q_B = 0.2$, then $x_{AB}^e = 0.11$. This value of x_{AB}^e is between $x_B = 0$ and $x_B = x_B^e$, as can be seen in Figure 9. In this case, suppose the adoption of B is very slow, i.e. it remains in the highlighted range for a long time. Once the fraction x_B crosses x_{AB}^e, all A users are incentivized to adopt AB converters. If the adoption of B is laggard and remains in the highlighted range for a long time, all current A users may adopt AB converters and become AB users, instead of switching to B. If this happens, we will require a higher fraction of BONLY adopters to incentivize AB users to switch to B. This, in turn, will slow down the adoption of B and can increase the cost of subsidizing the adoption of B. This resembles the current state of

IPv6 adoption today. Thus, in Case 2, we set the number of AONLY users to zero and study the system when all users are AB or BONLY.

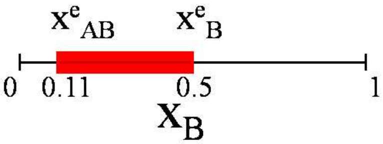

Figure 9: Range of x_B over which converters are actively adopted (highlighted) when $r_A = 0.1$ and $q_B = 0.2$.

When all users are either AB or BONLY, the equivalence point of a user is $x^e_B = \frac{S}{N\beta(2q_B - r_A)} - \frac{\alpha_B(1-t_B) - \alpha_A(1-r_A)}{2N\beta(2q_B - r_A)} + \frac{q_B - r_A}{2q_B - r_A}$. Ignoring S, and comparing with x^*_B from Section 4.1.2, we find that $x^*_B = x^e_B + \frac{\alpha_B(1-t_B) - \alpha_A(1-r_A)}{N\beta(2q_B - r_A)}$. Unlike in the case with no converters, we cannot identify an order relationship between x^*_B and x^e_B without plugging in the various parameter values. For example, if $t_B = 0$ and $r_A = 0$, then $x^*_B = x^e_B + \frac{\alpha_B - \alpha_A}{4N\beta q_B}$, and thus $x^*_B \geq x^e_B$. As the converter efficiency increases, gap between the ease of switching individually versus collectively widens.

4.2.3 Case 3: AONLY *and* BA *users*

Using analysis similar to that in Case 2, we study Case 3, where only BA converters are available. The inequality $U_{BA} \geq U_{Bonly} + S$, keeping the fraction x_A fixed, leads to the condition that $x_A \geq$ the threshold $\frac{S - \alpha_A t_A + \alpha_B r_B}{(1 - q_A + r_B)\beta N} + \frac{r_B}{1 - q_A + r_B}$. As before, this threshold is close to 0 and hence implies that until a very large fraction of A users adopt B (i.e., until the fraction x_A becomes small), B users will always have an incentive to adopt a BA converter. Hence, if BA converters are available, and the adoption of B is slow, then all B users will adopt BA converters and become BA users. We are thus justified in ignoring BONLY users in Case 3. The analysis of Case 3 yields results similar to that in Case 2: The gap between switching individually versus collectively widens as the converter efficiency increases.

5. SIMULATION STUDY

Our analysis in the previous section considered only simple scenarios that were easily tractable for mathematical study. Many aspects of our model, like the switching behavior of users in the presence of randomness, are difficult to study by mathematical analysis alone. We built a custom simulator to study the dynamics of our model in these complex scenarios. In addition to supporting the observations in the previous section, our simulation results reveal three insights:

1. The adoption of a new network architecture accelerates when users get the news about other users adopting the new architecture more quickly, except when converters are super efficient.
2. The adoption of a new network architecture may stall if network effects do not fall within an upper and lower bound determined by the current network conditions.
3. Increasing AB or BA converter efficiencies is detrimental to the adoption of B in some scenarios.

A user in our simulation study closely resembles the user described in Section 3. In addition to the standalone utility (α), the network effect parameter (β) and technology type (AONLY, BONLY, AB and BA), a user is associated with a switching cost, a limit on the maximum number of switches and a degree of randomness in switching. Randomness in switching captures a user's indecisiveness in choosing between competing network architectures offering similar benefits. This randomness in switching is defined by a Switch Threshold (ST) and a Switch Probability (SP). We initialize the simulator with a pool of users having different technology types. At each simulation time instant, the simulator iterates through all users in random order. Using the formulae from Section 3, each user calculates the difference in the utilities associated with different technology types and the sum of his current utility and switching cost. A user switches to the technology type offering the largest difference that is greater than the threshold ST. If none of the differences is greater than ST, the user decides to switch or not with probability SP. If the user does decide to switch, he randomly chooses one of the technology types for which the absolute value of the difference is less than ST. A user will not switch if he has already reached the maximum switch limit[3].

We consider two models by which the information about a user's switch spreads to other users. In the ENDOFITER model, other users know about a switch only at the beginning of the next time instant (iteration). In the INSTANT model, all users immediately know about the switch.

Due to the absence of real data, we assume the parameter values shown in Table 2 for our simulation study. When simulating different scenarios, we varied the relevant parameters. We refer to ST=100,SP=0.25 as *low randomness* and ST=500,SP=0.25 as *high randomness*. Unless mentioned otherwise, the ENDOFITER switching model is used. The number of users is limited to 10 million to keep simulation run-times tractable. Simulation results were averaged over 5 runs under different randomness seeds.

5.1 Standalone Utility

The analysis in Section 4.1 showed that, from a total system utility standpoint, increasing the standalone utility offered by B increases the feasibility of complete B adoption. Simulation results in this section show that this result holds even from the standpoint of users individually making the decision to adopt B or not. We find that a sharp threshold exists for $\frac{\alpha_B}{\alpha_A}$, above which complete B adoption takes place.

Figure 10 tracks the adoption of B for different values of $\frac{\alpha_B}{\alpha_A}$, in the absence and presence of randomness in switching. Let us first consider the cases where there is no randomness in switching. When the standalone utility offered

---

[3] In practice, the maximum switch limit is 1. However, different switch limits were considered in order to study their effects on system convergence.

Converter Properties		User Properties	
q_A	0.1	β	0.001
q_B	0.1	α_A	1000
r_A	0.0	α_B	1900
r_B	0.0	Switching Cost	Uniformly random between 0 and 1500
t_A	0.0	Switch Threshold (ST)	0 (no randomness)
t_B	0.0	Switch Probability (SP)	0
		Maximum Number of Switches	No limit
Initial Population Distribution			
A	9000000	B	1000000
AB	0	BA	0

Table 2: Parameter values common across simulations

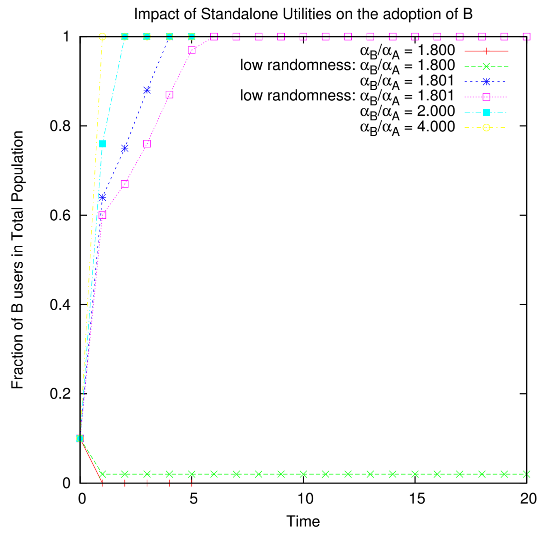

Figure 10: Penetration of B over time for different values of $\frac{\alpha_B}{\alpha_A}$ with and without randomness.

by B is less than or equal to 1.8 times that offered by A, all B users in the initial population switch to A due to the very high network benefits offered by the large initial population of A users. The penetration of B goes to 0 in a single simulation time instant. When $\frac{\alpha_B}{\alpha_A}$ is 1.801, complete adoption of B takes place in 4 simulation time units. The threshold of 1.8 for $\frac{\alpha_B}{\alpha_A}$ depends on the values chosen for our simulation parameters. Although the exact value of the threshold was discovered through simulations, analyzing the model equations in Section 3 and the switching cost distribution can provide a reasonable estimate. Adoption of B proceeds faster as $\frac{\alpha_B}{\alpha_A}$ increases – when $\frac{\alpha_B}{\alpha_A}$ = 4, complete adoption of B takes place in a single simulation time instant.

When the total utility (standalone utility + network effects) offered by B is comparable to that offered by A, after accounting for switching costs, there is a degree of randomness in a user's decision to switch or not. When $\frac{\alpha_B}{\alpha_A}$ is at or below the threshold, randomness in switching prevents the B penetration from dropping all the way to zero, as in the zero randomness case. When $\frac{\alpha_B}{\alpha_A}$ is greater than the threshold, randomness in switching slows down the complete adoption of B – for example, 6 simulation time instants versus 4 simulation time instants when $\frac{\alpha_B}{\alpha_A}$ = 1.801. This indicates that user indecisiveness in switching slightly aids B adoption when $\frac{\alpha_B}{\alpha_A}$ is below the threshold and is detrimental otherwise.

5.2 Network Effects

The total system utility analysis in Section 4.1 showed that increased network effects hinder the complete adoption of B. Simulation results in this section corroborate this result from the standpoint of individual users. The simulation results also indicate that complete adoption of B is impeded if network effects are too low.

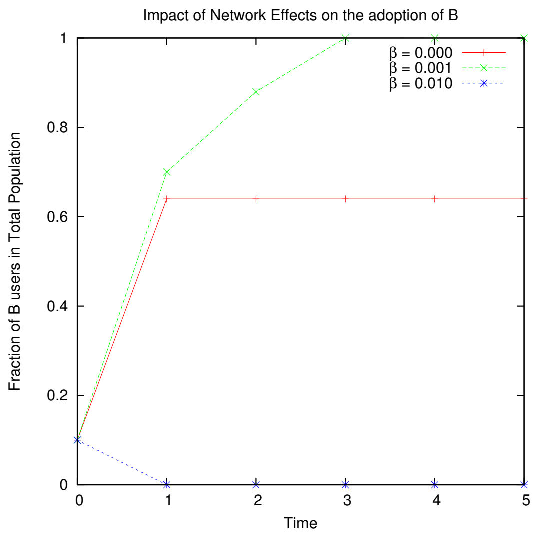

Figure 11: Penetration of B over time for different values of β without randomness in switching. $\frac{\alpha_B}{\alpha_A}$ = 1.9

Figure 11 shows the adoption of B for various values of the network effects importance parameter β. $\frac{\alpha_B}{\alpha_A}$ has been set to 1.9 to make switching to B attractive. The penetration of B stalls at 63% when network effects have no importance, i.e. $\beta = 0$. In the absence of the network benefits associated with the large number of A users, many users switch to B as it offers a higher standalone utility. However, some A users with high switching costs do not switch to B. In the absence of network effects, these A users do not have sufficient incentive to switch to B solely to enjoy the higher standalone utility offered by B, and adoption of B thus stalls.

The presence of network effects increases the total utility offered by B as more and more users switch to B, thus encouraging even more users to switch. For our chosen parameter values, this occurs when β = 0.001 and complete adoption of B takes place. However, if the network effects become too high, then adoption of B stalls. For example, when β = 0.01, the tremendous network effects associated with the large initial population of A users encourages the entire initial B population to switch to A in the very first simulation time instant. Thus network effects play an important role in the adoption of a new network architecture.

5.3 Converter Efficiencies

The analysis in Section 4.1 shows that, from a total utility standpoint, increasing AB and BA converter efficiencies reduces the barrier to B adoption. However, simulation results in this section show that this is not always true when B adoption is determined by switching decisions of individual users. Individual utility analysis in Section 4.2 introduced the complex system behavior caused by the presence of converters and showed that neither A nor B is a clear winner. Simulation results further illustrate this complexity and confirm the strong dependence on the parameter values chosen.

Simulation results indicate that increasing AB converter efficiency hinders the adoption of B. This is intuitively expected: Efficient AB converters reduce the total utility loss to be endured by the entire system in switching to B. However, from an individual user standpoint, highly efficient AB converters provide existing A users with a large portion of the network benefits associated with B and thus reduce their incentive to switch to B, with or without adopting BA converters. Thus, the availability of highly efficient AB converters is detrimental to the adoption of B.

On the other hand, the availability of very efficient BA converters should encourage A users to switch to B as they can continue to enjoy a large fraction of the network benefits associated with A users while reaping the higher standalone benefits of B. Simulation results indicate that increasing BA converter efficiency promotes B adoption in most scenarios. Figure 12 shows that at 92.5% BA converter efficiency, complete B adoption takes 2 simulation time instants longer than at 98% BA converter efficiency. Figure 12 also shows that increasing the BA converter efficiency hinders the adoption of B in some cases. When the converter efficiency is increased from 92.5% to 94.5%, adoption of B stalls at 95%. This surprising result can be explained as follows: At each simulation time instant, all A users with switching costs less than the extra utility offered by B switch from A to B. For the remaining A users to switch, the extra utility offered by B over A should increase as simulation time progresses. Higher BA converter efficiencies will increase the utility offered by B and encourage more users to switch to B. At the same time, since BA converters are two-way, higher BA converter efficiency will also increase the network benefits for the remaining A users and disincentivize them from switching to B. These conflicting effects determine how increasing BA converter efficiency influences B adoption.

Let us assume that BA converters cause no degradation in the benefits of B, i.e. $r_B = 0$, and hence all BONLY users adopt BA converters. On comparing the values of the utility difference $U_{BA} - U_{Aonly}$ at time instants 0 and 1, we find that the number of A → B switches should be greater than the following threshold in order to prevent the adoption of B from stalling:

$$N_{A->Bswitches} \geq \frac{N_{Bonly}}{2}(1/q_A - 1) \qquad (6)$$

Equation 6 shows that at high values of BA converter efficiency (i.e. lower q_A), a larger number of A → B switches are required to prevent the adoption of B from stalling. Although, higher converter efficiency leads to higher number of switches, complex interactions between other model parameters influence the number of switches and can make it fail to satisfy Equation 6. This happens when the BA converter efficiency is 94.5% under model parameter values used in our simulation. The main takeaway from this surprising result is that blindly increasing BA converter efficiency can sometimes prove detrimental to the complete adoption of B.

The results in this section indicate that increasing the efficiency of BA converters and keeping the efficiency of AB converters low promotes the adoption of B. One way this can be achieved is by letting BA users easily communicate with A users, while occasionally degrading the performance of the communication between AB and B users. Also, BA converters must be carefully engineered such that their efficiencies fall in the range that promotes complete B adoption.

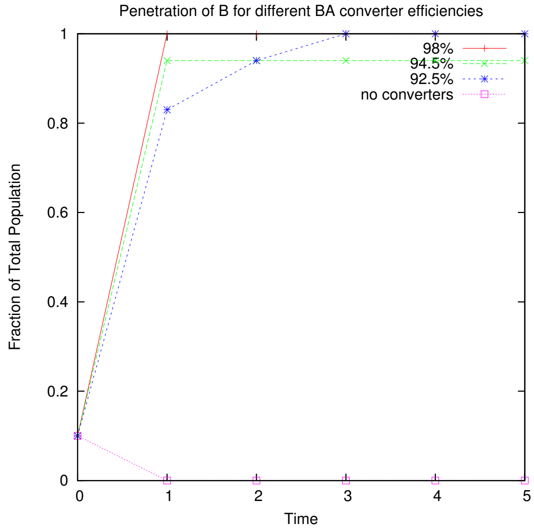

Figure 12: Penetration of B over time for different BA converter efficiencies when AB converters are 10% efficient.

5.4 Instant **versus** EndOfIter

Simulation results show that, in most scenarios, complete B adoption is faster under the Instant information spread model than under the EndOfIter model. More and more A users are encouraged to switch to B as news about the growing B population reaches them faster under Instant information spread. However, when BA converters are close to 100% efficient, Instant information spread hinders B adoption. This is because converters are two-way and Instant information spread quickly informs A users about the increased number of BA users they can now talk to, without switching to B. Instant information about increasing B penetration also encourages A users to directly switch to BOnly without getting stuck with BA converters they are loath to incur switching costs and discard even after all A users have switched to B. Hence, in most scenarios, new network architecture adoption will benefit if information about adoption is publicized and quickly dispersed to all users.

5.5 Performance hit of a BA converter

The analysis in Section 4.1 showed that adoption of B is infeasible from a total utility standpoint if the BA converter greatly degrades the benefits of B. In addition to supporting this conclusion from the standpoint of individual users, simulation results in this section indicate that a BA converter with some degradation encourages BA users to discard their converters and become BOnly users.

Figure 13 charts the adoption of B for different amounts of degradation caused by BA and AB converters. When neither converter causes any degradation, adoption of B quickly completes in 5 simulation time instants. As the degradation caused by the converters increases, adoption of B slows down because the benefits of adopting a converter decreases. Adoption of B stalls when the degradation caused by the converters reaches 5%. This result indicates that we should try to minimize the degradation caused by converters.

If the BA converter does not cause any degradation in the benefits of B, BA users have no incentive to incur switching costs and discard their BA converters even after all A users have disappeared. Although degradation slows down B adoption, a BA converter that slightly degrades the benefits of B may help in converting users to BOnly. Thus BA converters need to be engineered carefully with the right level of r_B in order to achieve a complete switch to BOnly users, without slowing down B adoption too much.

5.6 External benefits offered by a converter

The analysis in Section 4.1 showed that, from a total system utility standpoint, adoption of B is favored if the standalone benefits of A offered by a BA converter are increased and the standalone benefits of B offered by an AB converter are kept low. Simulation results in this section corroborate and quantify these results from the standpoint of individual users. When both AB and BA converters offer a small fraction (1%) of the standalone benefits of B and A respectively, complete adoption of B takes place. Complete adoption of B takes place as long as this fraction does not exceed 11.11%. The value of this threshold was estimated from the model equations in Section 3 and validated through simulations. B adoption stalls when the AB converter offers greater than 11.11% of the standalone benefits of B even when the BA converter offers the same fraction of the standalone benefits of A. Users of architecture A have no incentive to switch to B if most of B's standalone benefits are offered by an AB converter. If the BA converter does not offer any part of the standalone benefits of A, the threshold for the fraction of B standalone benefits provided by an AB converter, above which adoption of B stalls, is even lower – 5.26% for the simulation parameter values used in this study. These results urge us to increase the fraction of the standalone benefits of A offered by BA converters and decrease the fraction of the standalone benefits of B offered by AB converters.

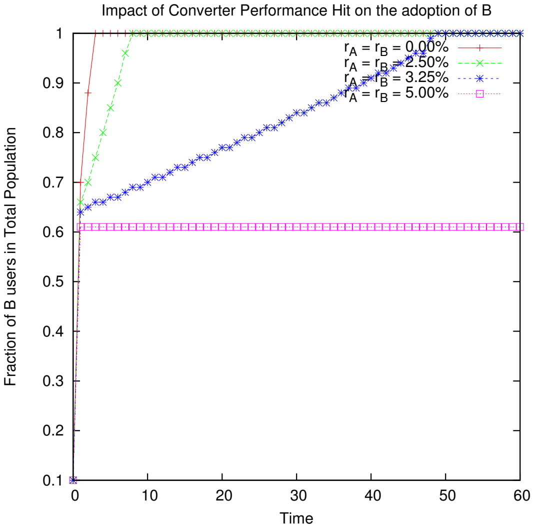

Figure 13: Penetration of B over time for different levels of converter performance hit.

6. CONTRIBUTIONS

The contributions of this paper are summarized below:

1. This paper reinforces our intuitive knowledge on new network architecture adoption through economic modeling, mathematical analysis and simulations.
2. Mathematical analysis and simulation studies help us in quantifying the relative importance of the various factors affecting the adoption of new network architectures. Our custom simulator can help study different new network architecture adoption scenarios.
3. This paper presents some unexpected results about new network architecture adoption. For example, increasing converter efficiency sometimes impedes new network architecture adoption.
4. This paper provides useful insights into designing converters that aid and not impede new network architec-

ture adoption. Our economic model and simulator can help analyze tradeoffs in converter design – for example, increasing converter efficiency versus decreasing degradation in benefits caused by converter operation.

5. The paper offers guidance on where industry or government efforts in promoting new network architecture adoption should focus. It is easier to encourage individual users to switch than to collectively coerce all users to adopt a new network architecture.
6. The paper shows that quickly spreading awareness of new network architecture adoption aids adoption, except when converters are very efficient.

7. LIMITATIONS

This paper has many limitations, some of which are being addressed by our ongoing work.

1. This paper is only a first step in modeling and analytically studying the adoption of new network architectures. Our model of new network architecture adoption is quite basic. Many missing aspects like interactions between ISPs, infrastructure vendors, application vendors and organizations, non-homogeneous switching costs and utility functions, economies of scale and learning as penetration increases and impact of geopolitical boundaries and influences are part of our ongoing research agenda.
2. The general nature of the model prevents it from capturing features which may be highly specific to a certain new network architecture and which may influence its adoption in ways other than that captured by our notion of user utility. Our basic model can be extended to capture these features.
3. A very big limitation of this paper is that the parameter values used in the analytical and simulation studies do not directly map onto real-world numbers. It is very difficult to accurately quantify the benefits provided by different network architectures and to measure the efficiency of converters. The lack of existing successful new network architecture adoptions results in the non-availability of real-world data to validate our model. We attempt only to draw general conclusions about the relative importance of the model parameters and how they relate to the real world.
4. Our simulation study is limited to 10 million users and does not consider the dynamic advent of new users. There is no direct mapping between simulation time and real-world clock time. The simulations compare different scenarios only based on virtual simulation time.

8. CONCLUSION

Studying the adoption of new network architectures through mathematical analysis and simulations is a fruitful research area. It corroborates many aspects of our intuitive understanding about new network architecture adoption, such as more superior the new network architecture is to the current one, the easier it is to deploy. At the same time, it brings our attention to unexpected results like higher efficiency converters do not always aid the adoption of a new network architecture. This paper increases our understanding of the various factors influencing adoption and gives insights into how we can improve converter design and engineering to hasten the adoption of a new network architecture.

9. REFERENCES

[1] Hexago. `http://www.hexago.com`.

[2] IPv6 Economic Impact Assessment - Final Report. National Institute of Standards and Technology, U.S. Dept. of Commerce, Oct 2005.

[3] R. Braden, T. Faber, and M. Handley. From Protocol Stack to Protocol Heap — Role-Based Architecture. In *HotNets 2002*.

[4] B. Briscoe, A. Odlyzko, and B. Tilly. Metcalfe's law is wrong - communications networks increase in value as they add members-but by how much? *IEEE Spectrum*, 43(7):34–39, July 2006.

[5] H. Chan, D. Dash, A. Perrig, and H. Zhang. Modeling adoptability of secure BGP protocol. In *SIGCOMM 2006*.

[6] J. P. Choi. The Provision of (Two-way) Converters in the Transition Process to a New Incompatible Technology. *Journal of Industrial Economics*, 45(2):139–153, 1997.

[7] D. Clark, R. Braden, A. Falk, and V. Pingali. FARA: reorganizing the addressing architecture. In *FDNA 2003*.

[8] J. Farrell and G. Saloner. Converters, Compatibility, and the Control of Interfaces. *The Journal of Industrial Economics*, 40(1):9–35, Mar 1992.

[9] P. Francis and R. Gummadi. IPNL: A NAT-extended internet architecture. In *SIGCOMM 2001*.

[10] R. Gold, P. Gunningberg, and C. Tschudin. A virtualized link layer with support for indirection. In *FDNA 2004*.

[11] A. Hovav, R. Patnayakuni, and D. Schuff. A model of Internet standards adoption: the case of IPv6. *Information Systems Journal*, 14, 2004.

[12] D. Joseph, J. Kannan, A. Kubota, K. Lakshminarayanan, I. Stoica, and K. Wehrle. OCALA: An Architecture for Supporting Legacy Applications over Overlays. In *NSDI 2006*.

[13] T. Koponen, M. Chawla, B. Chun, A. Ermolinskiy, K. Kim, S. Shenker, and I. Stoica. A Data-Oriented (and Beyond) Network Architecture. In *SIGCOMM 2007*.

[14] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana. Internet Indirection Infrastructure. In *SIGCOMM 2002*.

[15] X. Yang. NIRA: A new Internet routing architecture. In *FDNA 2003*.