Modeling the Adoption of new Network Architectures

Dilip Antony Joseph
John Chuang
Ion Stoica

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2007-41
http://www.eecs.berkeley.edu/Pubs/TechRpts/2007/EECS-2007-41.html

April 12, 2007

Modeling the Adoption of new Network Architectures

Dilip Joseph
dilip@cs.berkeley.edu

John Chuang
chuang@ischool.berkeley.edu

Ion Stoica
istoica@cs.berkeley.edu

Abstract

We propose an economic model based on user utility to study the adoption of new network architectures such as IPv6. We use analysis and simulation studies to understand the role of various factors such as user and network benefits, switching costs, and the existence of converters on new network architecture adoption. We find that carefully engineered converters, that offer a new network architecture user with partial benefits of the existing network architecture, hasten the adoption of the new network architecture.

1 Introduction

IPv6 has not achieved widespread adoption even after over a decade of existence. 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]. In the meanwhile, the networking research community has proposed many more new network architectures [3, 9], some of them radically different from IPv4 and some without clear transition and deployment mechanisms. What factors will aid the deployment of these new network architectures? In this report we construct a model of new network architecture adoption and analyze this question from an economics standpoint.

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 still 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.

The analytical and simulation studies confirm some of the obvious and intuitive observations about new network architecture adoption. For example, higher the standalone benefits offered by the new architecture, the faster is its adoption. Adoption of a new network architecture happens faster if users get the news about other users adopting the new architecture more quickly. Our study also exposes some non-obvious results and observations about new network architecture adoption. For example, adoption of a new network architecture stalls if the network effects do not fall within an upper and lower bound determined by the current network conditions. Another example is that increasing the efficiency of converters sometimes slows down the adoption of a new network architecture rather than quickening it. Some of the key insights revealed by the analysis and simulation study are as follows:

- New network architecture adoption needs to withstand a period of decreasing total system utility till a critical mass of users is reached. Incentives from the government or an industry champion can encourage users to adopt the new architecture and expedite the attainment of critical mass.

- Converters aid the adoption of new network architectures by reducing the loss in total utility before critical mass is attained. However, converters may be detrimental to complete adoption of the new network architecture unless they are carefully designed and engineered.

- Adoption of a new network architecture happens faster if users get the news about other users adopting the new architecture more quickly.

- New network architecture adoption stalls if the network effects do not fall within an upper and lower bound determined by the current network conditions.

We describe these and other results in detail in Sections 4 and 5.

This work is only an initial step in studying the adoption of new network architectures. Our 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 6.

2 Related Work

Adoption of new technologies and products has been extensively studied in Economics (For example, [6, 5]). 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 in the case of IPv6, there are no opposing organizations, each pushing its own technology. 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 [7] 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, [7] uses case studies to construct and illustrate the model and does not perform an analytical or simulation study. Unlike [2] and [7], our report focuses on the role of converters in the adoption of new network architectures. In [4], the authors model and simulate the adoption of secure BGP protocols and define the switching threshold as an adoptibility metric. Our report 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 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: (i) 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 host address configuration provided by IPv6. (ii) Network benefits derived from the ability to communicate with other users of the same architecture.

For example, *i3* [9] 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 report.

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 [8] 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 A users 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:

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

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 effects [6], with a single parameter β controlling the importance of the network effects. $\beta N x_{BA} x_B (1 - q_A)$ 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.

Similar to Equation 1, the utility enjoyed by a B user who does not use an BA converter, i.e. a BONLY user, is given by:

$$U_{Bonly} = \alpha_B + \beta N x_B + \beta N x_{AB} x_A (1 - q_B) \quad (2)$$

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

$$U_{AB} = (1 - r_A)(\alpha_A + \beta N x_A) + \beta N x_B (1 - q_B) + t_B \alpha_B \quad (3)$$

Parameter r_A captures the potential degradation caused by an AB converter on the utility offered by A. For example, 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.

For simplicity, we treat all B users (both BONLY and BA) alike and apply only a single degradation factor $(1 - q_B)$ 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 fraction (t_B) of the standalone benefits offered by B. For example, the OCALA proxy provides *i3*-style mobility support to a user's applications, while enabling communication with other *i3* users.

Similar to Equation 3, the utility enjoyed by a B user who uses a BA converter, i.e. a BA user, is given by:

$$U_{BA} = (1 - r_B)(\alpha_B + \beta N x_B) + \beta N x_A (1 - q_A) + t_A \alpha_A \quad (4)$$

The total utility enjoyed by all the users in the system is given by the following expression:

$$\begin{aligned} TU = {} & N(1 - x_{AB}) x_A U_{Aonly} + N x_{AB} x_A U_{AB} + \\ & N(1 - x_{BA}) x_B U_{Bonly} + N x_{BA} x_B U_{BA} \end{aligned} \quad (5)$$

4 Mathematical Analysis

In this section, we analyze the model formulated in the previous section to quantify the impact of various parameters on the adoption of the new network architecture B. We first study the case when all users

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

Table 1: Notations used in this report.

collectively make the switching decision in order to maximize the total utility of all users in the system. We then study the case where a user makes the switching decision purely to maximize his individual utility.

4.1 Analysis of the 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[1]. 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 4 distinct cases: (0) AONLY users, (2) AONLY and BONLY users, (2) AB and BONLY users, and (3) BA and AONLY users.

Our analysis leads us to two main observations. First, as expected, there is a period of decreasing total system utility associated with the adoption of a new network architecture. Government intervention and economic incentives may be needed to achieve a critical penetration level. Second, while converters aid the adoption of new network architectures in the initial phase, they may be detrimental to complete adoption if they are "too good". This observation advocates a controversial strategy of intentionally keeping the standalone utility of the converter low to promote complete adoption of the new architecture.

4.1.1 Case 0: AONLY users

In Case 0, we assume that everyone in the system uses A alone, i.e. $x_A = 1$, $x_{AB} = x_{BA} = x_B = 0$. The total utility of the system is given by:

$$TU_0 = N\alpha_A + N^2\beta \tag{6}$$

Unsurprisingly, the total utility increases with increasing N.

4.1.2 Case 1: AONLY and BONLY users

When no AB and BA converters are available, all users in the system are of type AONLY or BONLY, i.e. $x_{AB} = x_{BA} = 0$. We study the total utility under different system conditions as the fraction of B users varies.

The total utility is at its minimum when the fraction of B users, x_B, equals $\frac{1}{2} + \frac{\alpha_A - \alpha_B}{4N\beta}$. If technologies A and B have equal standalone utilities, i.e. $\alpha_A = \alpha_B$, then the minimum utility occurs when there are equal number of A and B users.

---

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

In the initial phase of B adoption (when x_B is low), adoption is infeasible from a total utility standpoint—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 initial switching 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 achieving 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 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 higher the standalone benefit offered by B, easier is its adoption.

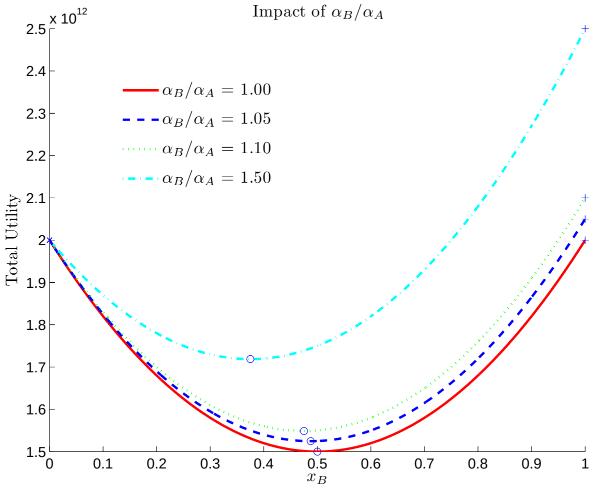

Figure 1: Total utility for different standalone utilities of B

Figure 2 shows the impact of network effects in Case 1. As $N\beta$ increases, the total utility increases and the fraction of B users to overcome the trough increases. The size of the trough also increases - both in terms of absolute utility and as a fraction of the utility at $x_B = 0$. This implies that technology B has a greater chance of adoption when the network effects are lower.

Figure 2 also shows that the number of users has a greater effect on the total utility than the network effect parameter β. In terms of absolute utility, it appears to be more difficult to overcome the trough when there are more number of users than when β is higher. The increased utility obtained by completely switching to B is also higher. However, if measured as a percentage of the utility at $x_B = 0$, the size of the trough and the increased benefits are identical when $N = 2e + 9, \beta = 1e - 6$ and when $N = 1e + 9, \beta = 2e - 6$.

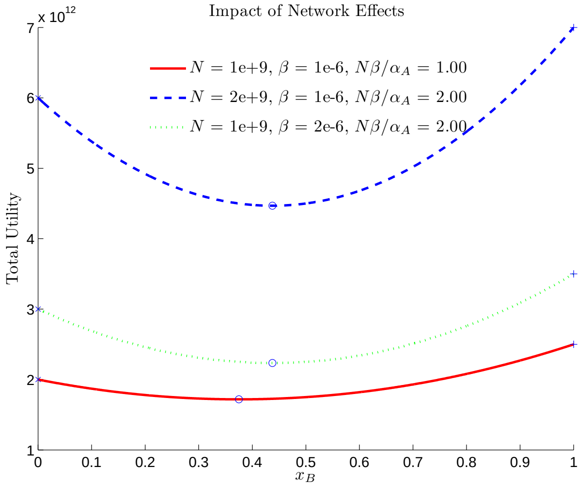

Figure 2: Impact of Network Effects ($\alpha_B/\alpha_A = 1.5$)

4.1.3 Case 2 : AB and BONLY users

In Case 2, we study the system utility when all A users have deployed AB converters and none of the B users have BA converters (possibly due to non-availability), i.e. $x_{AB} = 1$ and $x_{BA} = 0$. If $U_{AB} \geq U_{Aonly}$, there is clear incentive for all A users to deploy AB converters. The total system utility is minimized (if $q_B > r_A/2$) or maximized (if $q_B < r_A/2$) when

$$x_B^* = \frac{q_B - r_A}{2q_B - r_A} + \frac{\alpha_A(1 - r_A) - \alpha_B(1 - t_B)}{2\beta N(2q_B - r_A)} \quad (7)$$

$U_{AB} \geq U_{Aonly}$ holds only if $x_B \geq \frac{r_A(\alpha_A+\beta N)-t_B\alpha_B}{\beta N(1-q_B+r_A)}$. If the AB converter causes no degradation in the standalone and network benefits of A, i.e. $r_A = 0$, then this condition trivially holds true for all values of $0 \leq x_B \leq 1$.

Let us assume that technology B offers 1.2 times the standalone benefit of A and the AB converter does not cause any degradation in the standalone or network benefits associated with A nor does it provide A with any of the standalone benefits associated with technology B. Thus $\alpha_B/\alpha_A = 1.2$, $r_A = 0$ and $t_B = 0$.

Figure 3 shows the impact of the efficiency (q_B) of the AB converter. As the converter becomes more efficient ($q_B \rightarrow 0$), the trough in total utility needed to be overcome before all the users convert to technology B decreases. The fraction of early adopters of B required to get across the trough also decreases as the efficiency of the AB converter increases. In Figure 3, the system can smoothly move from $x_B = 0$ to $x_B = 1$ without any trough if the AB converter has efficiency greater than 90%. Thus, from a total utility standpoint, more efficient AB converters help in complete adoption of B.

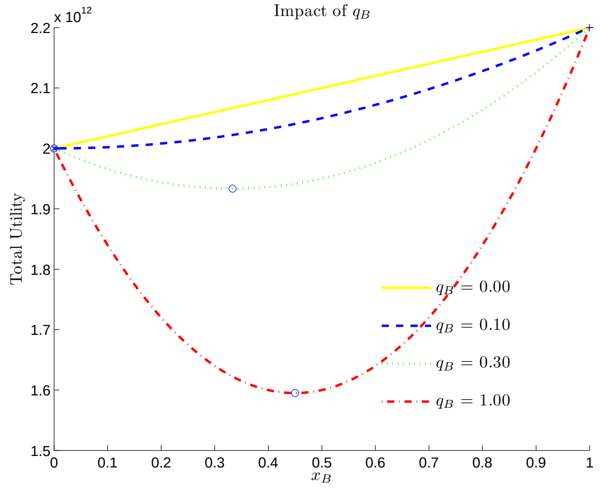

Figure 3: Impact of AB Converter Efficiency ($N = 10^9, \beta = 10^{-6}, \alpha_B/\alpha_A = 1.2$)

What happens if an AB converter, in addition to enabling A users to communicate with B users, also provides an A user with some part of the standalone benefits associated with B? Equation 7 implies that as t_B increases, a larger number of early adopters of B is required for the system to overcome the decrease in utility and move towards full B deployment. In Figure 4, as t_B increases, full deployment of B becomes more difficult. The fraction of early adopters of B required to overcome the trough increases. For sufficiently large values of t_B, full deployment of B looks impossible. If an AB converter provides a sufficiently large portion of the standalone benefits of B, then there is no incentive for a user to switch to technology B.

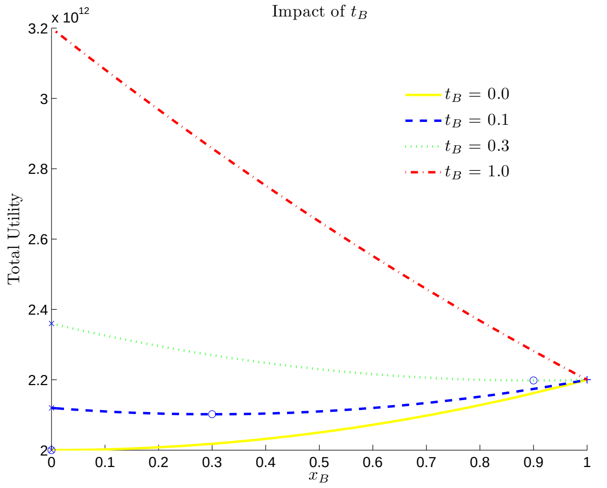

Figure 4: Impact of t_B ($N = 10^9, \beta = 10^{-6}, \alpha_B/\alpha_A = 1.2, q_B = 0.1, r_A = 0$

When $r_A > 2q_B$, the AB converter imposes a heavy degradation on the standalone and network benefits associated with A. Even in this scenario, as Figure 5 shows, a high t_B does not aid in the complete deployment of B. It increases the total utility at $x_B = 0$ and hence reduces the additional utility to be gained by switching to $x_B = 1$. This leads to a controversial question - In order to promote the adoption of the new technology B, should we on purpose ensure that t_B is low?

4.1.4 Case 3 : BA and AONLY users

In Case 3, we study the system utility when all B users have deployed BA converters and none of the A users have deployed AB converters, i.e. $x_{BA} = 1$ and $x_{AB} = 0$. If $U_{BA} \geq U_{Bonly}$, there is clear incentive for all B users to deploy BA converters.

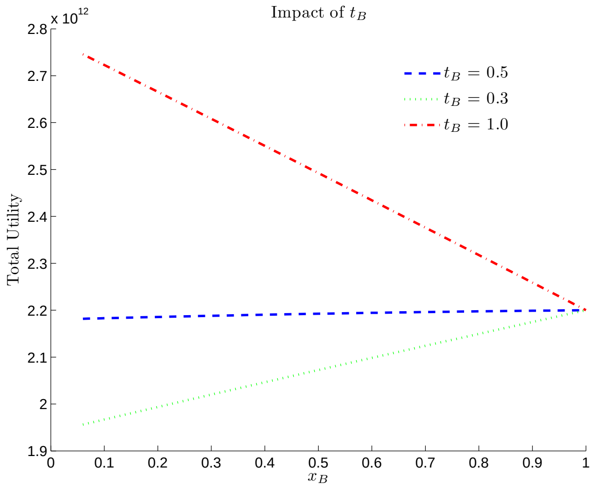

Figure 5: Impact of t_B ($N = 10^9, \beta = 10^{-6}, \alpha_B/\alpha_A = 1.2, q_B = 0.1, r_A = 0.21$). We consider only the range $0.06 \geq x_B \leq 1$ in order to ensure that $U_{AB} \leq U_A$ holds.

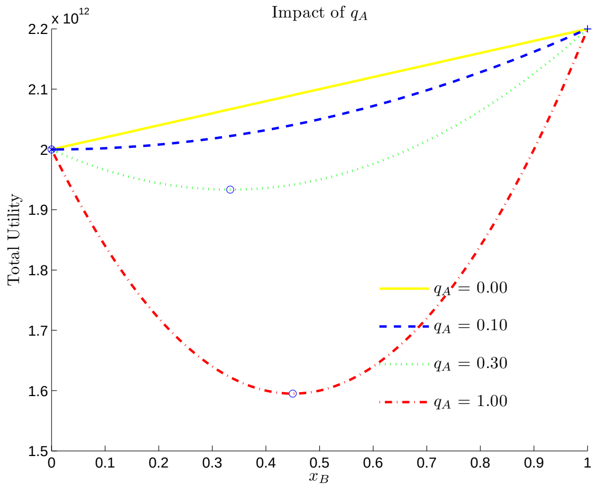

Figure 6: Impact of BA Converter Efficiency ($N = 10^9, \beta = 10^{-6}$, $\frac{\alpha_B}{\alpha_A} = 1.2$)

The total system utility is minimized (if $q_A > r_B/2$) or maximized (if $q_A < r_B/2$) when

$$x_B^* = \frac{q_A}{2q_A - r_B} + \frac{\alpha_A(1-t_A) - \alpha_B(1-r_B)}{2\beta N(2q_A - r_B)} \quad (8)$$

$U_{BA} \geq U_{Bonly}$ holds only when $x_B \leq (\alpha_A t_A + N\beta(1-q_A))/(N\beta(1+r_B-q_A))$. For $r_B = 0$ and $t_A = 0$, the condition reduces to $x_B \leq 1$, which is true by definition.

Figure 6 shows the impact of the degradation of the BA converter under conditions similar to Figure 3 ($\alpha_B/\alpha_A = 1.2$, $r_B = 0$ and $t_A = 0$), which analyzed the impact of an AB converter. The two figures exhibit identical trends. This result directly follows from Equation 5, on substituting appropriate parameter values. Thus, whether we are building an AB converter or a BA converter, greater converter efficiency increases the widespread deployment chances of technology B.

Figure 7 shows that increasing t_A increases the total additional system utility attained as more and more users adopt technology B. Higher values of t_A aid the widespread deployment of technology B. In Case 2, we saw that higher values of t_B hamper deployment of B. Thus, in order to hasten the deployment of technology B, we should build BA converters that offer a substantial portion of the standalone benefits of A, or build AB converters which do not offer the standalone benefits of B.

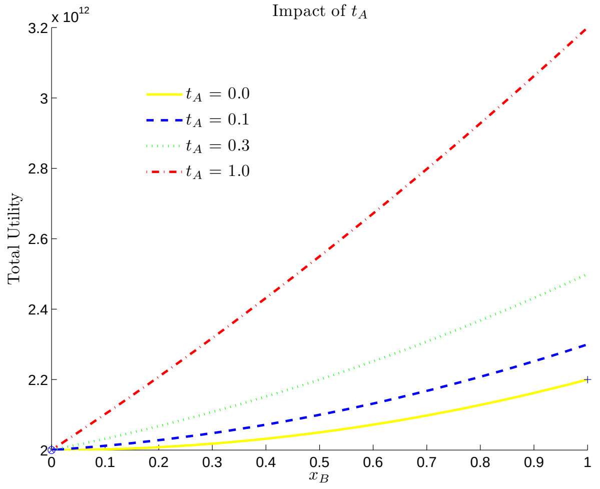

Figure 7: Impact of t_A ($N = 10^9, \beta = 10^{-6}$, $\frac{\alpha_B}{\alpha_A} = 1.2, q_A = 0.1, r_B = 0$

Figure 8 shows that it is important to minimize any degradation caused by a BA converter in the standalone and network benefits associated with B. Full deployment of B is viable from a total system utility standpoint as long as r_B is smaller than a threshold.

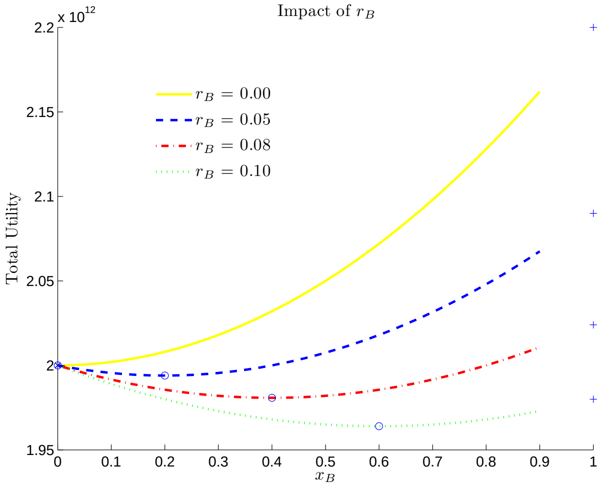

Figure 8: Impact of r_B ($N = 10^9, \beta = 10^{-6}$, $\frac{\alpha_B}{\alpha_A} = 1.2$, $q_A = 0.1$). We consider only the range $x_B \leq 0.9$ in order to ensure that $U_{BA} \leq U_B$ holds.

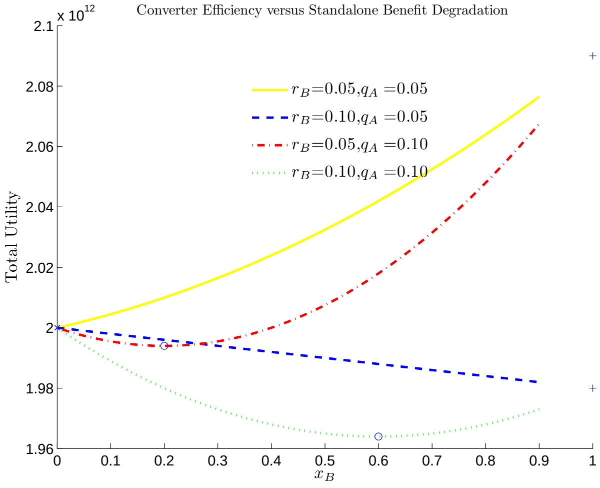

Figure 9: Impact of r_B ($N = 10^9, \beta = 10^{-6}$, $\frac{\alpha_B}{\alpha_A} = 1.2$). We consider only the range $x_B \leq 0.9$ in order to ensure that $U_{BA} \leq U_B$ holds.

4.2 Analysis of Individual Utilities

The decision to adopt a new technology is often made by users themselves, while considering only their individual benefits. In this section, we compare the individual utilities associated with the different technology adoption scenarios and study the factors which promote switching to a different technology. The main result of this section is that incentivizing individual users to switch to a new architecture is easier than coercing users to switch collectively. Furthermore, we show that as converter efficiency increases, the gap between the ease of switching individually and collectively widens. As in Section 4.1, for simplicity, we limit our analytical study to three distinct cases.

4.2.1 Case 1 : AONLY and BONLY users

In Case 1, we assume that users have deployed either technology A or technology B but with no converters, i.e. $x_{AB} = x_{BA} = 0$. At a particular instant of time, a user switches from A to B if $U_{Bonly} \geq U_{Aonly} + S$, where S is the switching cost. This holds true if $x_B \geq \frac{1}{2} + \frac{S-(\alpha_B-\alpha_A)}{2N\beta}$. When $x_B = \frac{1}{2} + \frac{S-(\alpha_B-\alpha_A)}{2N\beta}$, the user is ambivalent between technologies B and A. We call this value of x_B the *Equivalence Point*, x_B^e. The equivalence point is lowered by a higher value of α_B and by lower network effects (assuming $\alpha_B \geq \alpha_A$) – lesser number of existing B users are required to encourage more users to switch to B. If we assume a zero switching cost, the expression for x_B^e is similar to the expression for the point of least total utility (Section 4.1.2), $x_B^* = \frac{1}{2} - \frac{\alpha_B-\alpha_A}{4N\beta}$. Comparing, x_B^* and x_B^e, we find that $x_B^* = x_B^e + \frac{\alpha_B-\alpha_A}{4N\beta}$. Assuming $\alpha_B \geq \alpha_A$, this implies that we need a lower seed population of B alone users to entice other users to switch to B individually rather than collectively. Thus it may be more rewarding to focus on getting individual users to switch to B rather than trying to switch the whole population in one go. This becomes more prominent if technology B is very superior to A, as the gap between x_B^* and x_B^e widens. 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.

4.2.2 Case 2 : AB and BONLY users

Now we consider the case when all A users have converters, i.e. $x_{AB} = 1$, and no B users have converters, i.e. $x_{BA} = 0$. This makes sense only if $U_{AB} \geq U_A$, i.e. $x_B \geq \frac{\alpha_A r_A - \alpha_B t_B - N\beta r_A}{(1-q_B+r_A)N\beta}$. This trivially holds if $r_A = 0$.

A user with an AB converter switches to B if $U_B \geq U_{AB} + S$. The equivalence point $x_B^e = \frac{S}{N\beta(2q_B - r_A)} - \frac{\alpha_B(1-t_B) - \alpha_A(1-r_A)}{2N\beta(2q_B - r_A)} + \frac{q_B - r_A}{2q_B - r_A}$. Comparing with x_B^* from Section 4.1.3, we find that $x_B^* = x_B^e + \frac{\alpha_B(1-t_B) - \alpha_A(1-r_A)}{N\beta(2q_B - r_A)}$. Unlike in the case with no converters, we cannot identify an order relationship between x_B^* and x_B^e without plugging in the various parameter values. For example, if $t_B = 0$ and $r_A = 0$, then $x_B^* = x_B^e + \frac{\alpha_B - \alpha_A}{4N\beta q_B}$, $x_B^* \geq x_B^e$. As the converter efficiency increases, gap between the ease of switching individually and collectively widens.

4.2.3 Case 3 : BA and AONLY users

Let us now consider the case when all B users have converters and no A users have converters, i.e. $x_{BA} = 1$ and $x_{AB} = 1$. One reason why $x_{AB} = 0$ could be due to the non-availability of an AB converter. $x_{BA} = 1$ makes sense only if $U_{BA} \geq U_B$. This trivially holds true if $r_B = 0$.

An A user switches to B with a BA converter if $U_{BA} \geq U_A$. The equivalence point is $x_B^e = \frac{q_A}{2q_A - r_B} - \frac{\alpha_B(1-r_B) - \alpha_A(1-t_A)}{N\beta(2q_A - r_B)}$. Comparing with x_B^* from Section 4.1.4, $x_B^* = x_B^e + \frac{\alpha_B - \alpha_A}{4N\beta q_A}$. As in case 2, whether x_B^e is greater or whether x_B^* is greater depends on the parameter values.

4.3 Take Aways

- Analysis of the total utility of the entire population shows a trough in the total utility that needs to be overcome for complete adoption of B. A critical mass of early adopters of B is required to go past the point of minimum total utility in this trough. Once past the minimum point, system dynamics to maximize the total utility lead the entire population to adopt B.

- Increasing the standalone utility of B (α_B) decreases the depth of the trough as well as decreases the number of early adopters of B required to go past the trough in total utility.

- Higher network effects ($N\beta$) makes adoption of B more difficult – it increases the depth of the trough as well as increases the number of of early adopters of B required to go past the trough in total utility. The number of users (N) has a higher effect on the absolute magnitude of the trough depth than β. Both N and β have the same influence on the trough depth, when considered as a percentage of the total utility when all users are AONLY.

- Both AB and BA converters aid the adoption of B by decreasing the trough as well as by reducing the critical mass of BONLY users required in the initial population. More efficient customers are more effective in aiding the deployment of B.

- For speedy adoption of B, an AB converter should not provide a large portion of the standalone benefits of B, even in order to offset any degradation in the standalone benefits of A caused by using the AB converter. Providing a portion of the standalone benefits of A in a BA converter aids the deployment of B.

- Complete adoption of B appears feasible from a total utility standpoint only if the self-degradation caused by the use of converters is below a threshold.

- The fraction of B users in the total population at which two different user types offer the same utility to an individual user is called the equivalence point for those two user types. In the absence of converters, the equivalence point of AONLY and BONLY is always less than the minimum point in the total utility curve. This means that it is easier to convince each individual user to adopt B than trying to collectively move the entire system past the minimum point in the total utility curve, i.e. a smaller number of of early adopters of B is required. There is no order relationship between the equivalence point and the minimum point of the trough when converters are present. The gap between the equivalence point and the minimum point widens with increasing converter efficiency. The gap also widens with increasing α_B irrespective of the presence or absence of converters.

5 Simulation Study

Some aspects of our model, like the switching behavior of users in the presence of randomness, are difficult to study by mathematical analysis alone. We use a custom simulator to study the behavior of our model in these complex scenarios. In addition to supporting the observations in the previous section, our simulation results reveal two key insights. First, 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. Second, the adoption of a new network architecture may stall, if the network effects do not fall within an upper and lower bound determined by the current network conditions.

Each user in our simulation study closely resembles the user model described in Section 3. In addition to the standalone utility (α), the network effect parameter (β) and technology type (AONLY, BONLY, AB and BA), each user is associated with a switching cost, a limit on the maximum of number of switches and a degree of randomness in switching. Randomness in switching is defined by a Random Switch Threshold (RST) and a Random Switch Probability (RSP). We initialize the simulator with a pool of users having different technology types. At each instant of simulation time, 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. The user switches to the technology type offering the largest difference which is greater than the RST. If none of the differences is greater than the RST, the user decides to switch or not with probability RSP. 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 the RST. A user will not switch if he has already reached the maximum switch limit. 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.

Table 2 shows the parameter values used in the simulations. When simulating different scenarios, we varied the relevant parameters. We refer to RST=100,RSP=0.25 as *low randomness* and RST=500,RSP=0.25 as *high randomness*. Unless explicitly mentioned otherwise, the ENDOFITER switching model is used. The number of users is limited to 10 million to keep simulation run-times tractable. We admit that the absolute values chosen for most of the parameters have no direct bearing to the real world. The observations which we summarize below focuses on the relative importance of different parameters. Real-life parameter values, if available in the future, can be easily plugged into our simulator.

We start by analyzing the importance of the standalone utilities offered by A and B in the next section. Many interesting behaviors and observations are common across analysis sections. In order to avoid repetition, these are explained in detail only in the first section. Hence, the first section is much larger in size than the rest.

5.1 Standalone Utilities

We study the importance of the standalone utilities (α_A and α_B) of technologies A and B by varying the $\frac{\alpha_B}{\alpha_A}$ ratio. Adoption of B lags till $\frac{\alpha_B}{\alpha_A}$ reaches a critical threshold, after which users switch to BONLY or BA rapidly. We expect the new technology B to be superior to A and hence the ratio $\frac{\alpha_B}{\alpha_A}$ to be greater than 1. Nevertheless, we start with case when both A and B have the same standalone utility.

When $\frac{\alpha_B}{\alpha_A}$ = 1, there is no incentive for any of the users to switch to B or to adopt a BA converter. This is due to the tremendous network effects associated with the large number of AONLY users in the initial population. Under the ENDOFITER model (Figure 10), the BONLY users in the initial population immediately switch to AB while some AONLY users

Converter Properties		User Properties	
q_A	0.1	β	0.001
q_B	0.1	α_A	1000
r_A	0.0	α_B	1800
r_B	0.0	Switching Cost	Uniformly random between 0 and 1500
t_A	0.0	Random Switch Threshold (RST)	∞
t_B	0.0	Random Switch Probability (RSP)	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

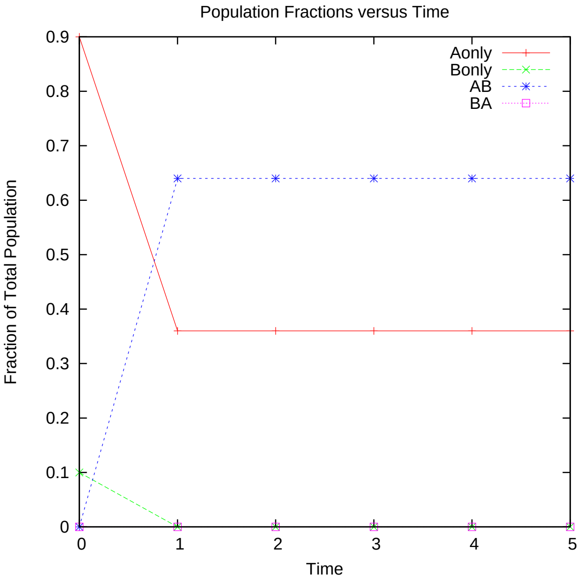

Figure 10: Population distribution when $\frac{\alpha_B}{\alpha_A} = 1$ under the ENDOFITER information spread model, no randomness and no switch limit.

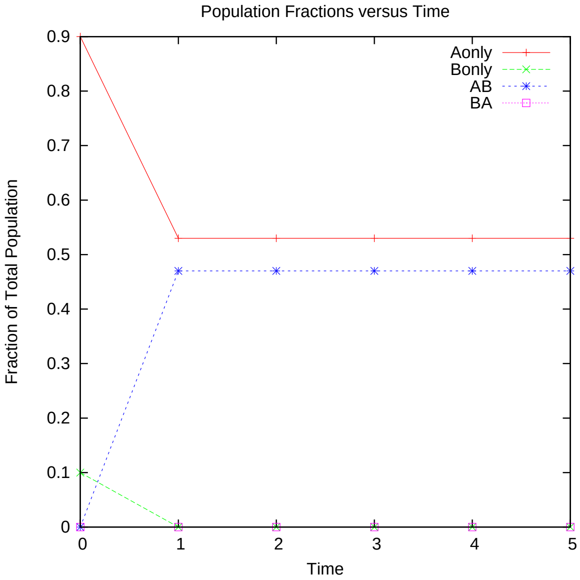

Figure 11: Population distribution when $\frac{\alpha_B}{\alpha_A} = 1$ under the INSTANT information spread model, no randomness and no switch limit.

with low switching costs adopt an AB converter. However, under the INSTANT model (Figure 11), most users become AONLY and do not adopt an AB converter. This occurs because the news about BONLY users switching to AB immediately reaches all users contemplating adoption of an AB converter for communicating with BONLY users. Lesser the number of BONLY users, lesser is the necessity of an AB converter. In either model, all BONLY and BA users disappear after the very first time instant.

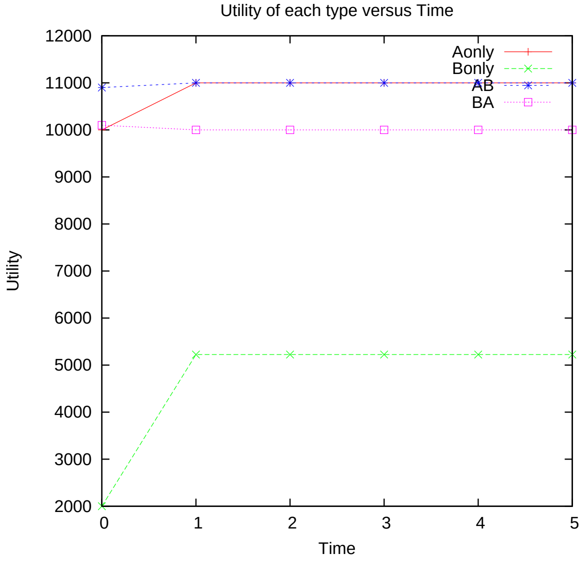

Figure 12: Individual utilities of the various technology types when $\frac{\alpha_B}{\alpha_A} = 1$ under the INSTANT information spread model, no randomness and no switch limit.

After the number of BONLY and BA users in the system goes to 0, AONLY and AB users have identical utilities (Figure 12), as the AB converter does not degrade the benefits associated with A (i.e. $r_A = 0$). This means that an AB user has no incentive to incur switching costs and give up his AB converter, even if the converter is useless when there are no BONLY users. In the ENDOFITER model, the number of AB users at convergence is greater than the number of AONLY users. However, under high randomness, the number of AONLY users is greater than the number of AB users. High randomness encourages users to jump from AB to AONLY even if the utilities offered by AONLY and AB are identical. In the INSTANT model, the number of AONLY users is greater than the number of AB users irrespective of randomness.

When randomness in switching is present, if no limit on the number of switches per user is imposed, users forever keep switching back and forth between AONLY and AB (Figure 13) as the utilities offered by AONLY and AB are identical after all B users have adopted A. However, even under high randomness, the number of AONLY and AB users converges after a few iterations. If we limit the maximum number of switches per user to 2 , in the ENDOFITER model, there are more AONLY users than AB users at convergence. If we limit the number of switches to 1, the system converges to having more AB users than AONLY users. This is because users cannot discard the AB converters they initially adopted when they were unaware that most of the BONLY users had switched to AONLY or AB. This does not occur in the INSTANT model – the number of AONLY users is always greater than the number of AB users at convergence, irrespective of the maximum number of switches allowed. The maximum switch limit does influence the magnitude of the difference between the number of AONLY and AB users at convergence in the INSTANT model. If only one switch is allowed (Figure 14), the number of AONLY users is only slightly greater than the number of AB users – the users who initially converted to AB cannot discard their converters even if they wanted to. If two switches are allowed (Figure 15), the number of AONLY users is much greater than the number of AB users, as most AB users discard their converters on finding that all BONLY users have switched to AONLY or AB. The number of AONLY users in this case is even greater than the case when infinite number of switches are allowed (Figure 16). This is because some AONLY users randomly switch to AB as part of the continuous back and forth switching between AONLY and AB. The maximum switch limits do not affect the scenarios where there is no randomness.

Complete adoption of B never happened in any of the scenarios were $\frac{\alpha_B}{\alpha_A} = 1$. In real life, the new technology B will have higher standalone utility than A. For example, IPv6 has more number of IP addresses than IPv4 and also enables host auto-configuration. Thus, $\frac{\alpha_B}{\alpha_A}$ is greater than 1. There exists a sharp threshold for $\frac{\alpha_B}{\alpha_A}$ above which complete adoption of B takes place. The value of this threshold can be analytically derived from the equations in Section 3. For the model parameters chosen in our simulation, this threshold is 1.8. Under zero or low randomness, the system converges to a combination of AONLY and AB users if α_B is 1.8 times α_A. However, under high randomness, the number of AONLY and AB

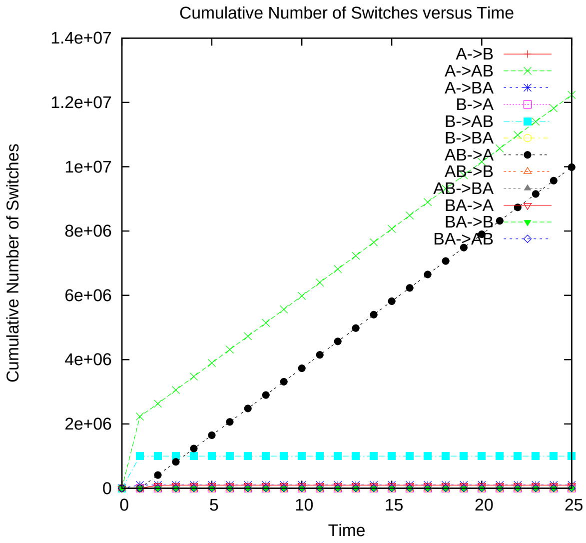

Figure 13: Cumulative number of switches between the various technology types $\frac{\alpha_B}{\alpha_A} = 1$ under the INSTANT information spread model, high randomness and no switch limit.

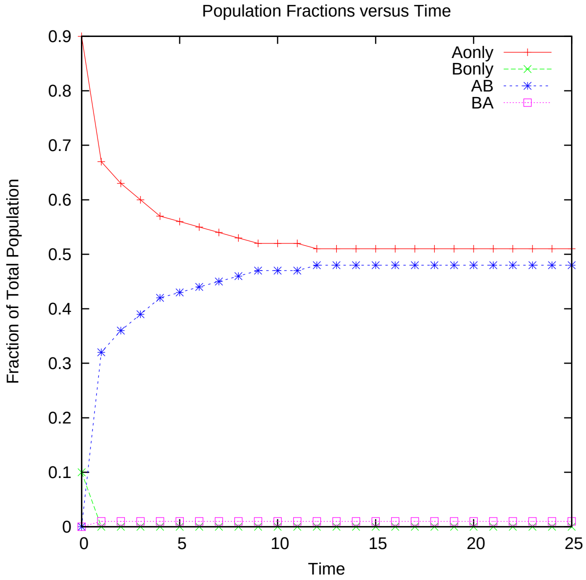

Figure 14: Population distribution when $\frac{\alpha_B}{\alpha_A} = 1$ under the INSTANT information spread model, high randomness and 1 switch limit.

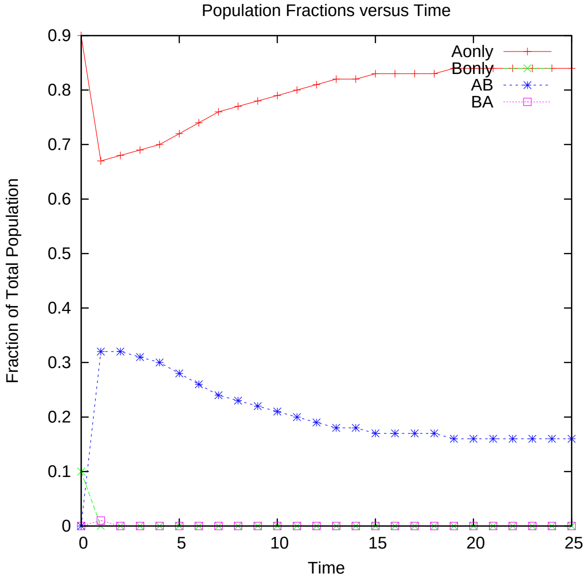

Figure 15: Population distribution when $\frac{\alpha_B}{\alpha_A} = 1$ under the INSTANT information spread model, high randomness and 2 switch limit.

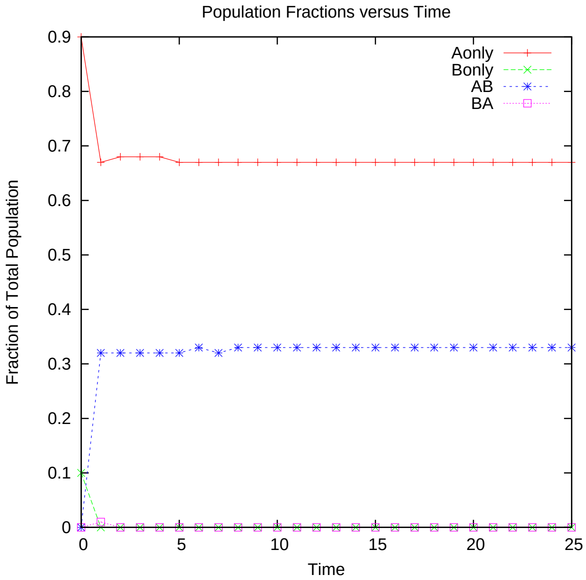

Figure 16: Population distribution when $\frac{\alpha_B}{\alpha_A} = 1$ under the INSTANT information spread model, high randomness and no switch limit.

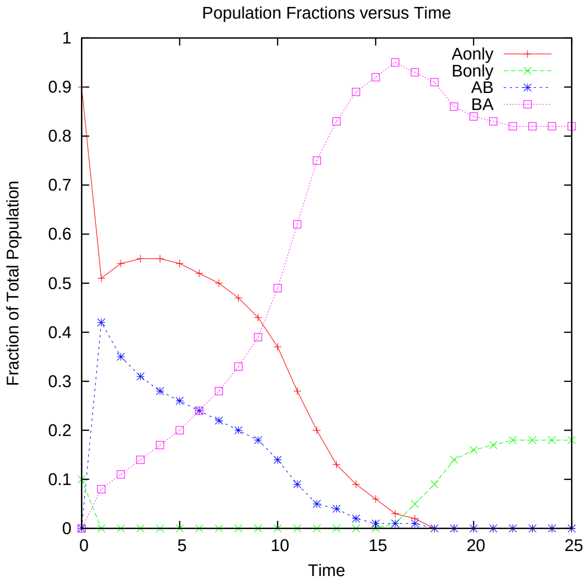

Figure 17: Population distribution when $\frac{\alpha_B}{\alpha_A} = 1.8$ under the EndOfIter information spread model, high randomness and no switch limit.

users in the system slowly declines till all users become Bonly or BA (Figure 17). Complete adoption of B is faster in the Instant model (7 iterations) (Figure 18) than in the EndOfIter model (16 iterations) – more and more users adopt B quickly if the news about other users' adoption of B reaches their ears quickly. This implies that publicizing adoption statistics is very important. In the Instant model, adoption of B stalls if we limit each user to a single switch (Figure 19). This is obvious as all users initially switch to Aonly or AB due to lack of information about the ongoing adoption of B, and get stuck with their initial choice.

When $\frac{\alpha_B}{\alpha_A}$ is greater than the threshold of 1.8 adoption of B happens even in the EndOfIter model, irrespective of randomness (Figure 20). As before, adoption of B is faster in the Instant model (2 iterations) (Figure 21) than in the EndOfIter model (4 iterations). However, contrary to prior behavior where randomness aided the adoption of B, randomness in switching slows down the adoption of B - 3 iterations versus 2 (Figure 22). Thus randomness in switching is beneficial to the adoption of B when $\frac{\alpha_B}{\alpha_A}$ is below the threshold and detrimental otherwise.

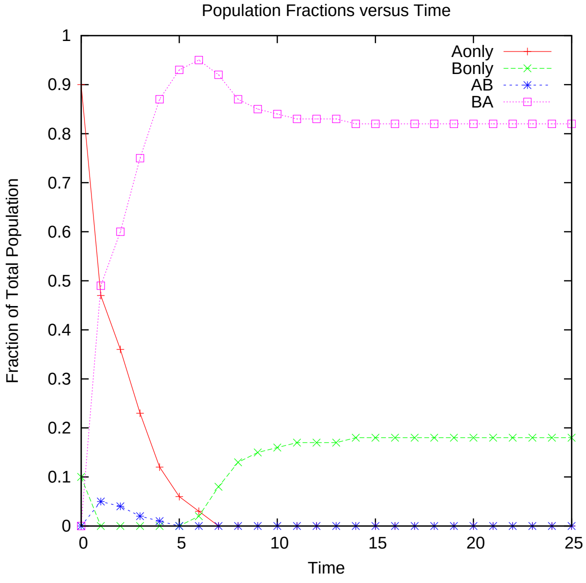

Figure 18: Population distribution when $\frac{\alpha_B}{\alpha_A} = 1.8$ under the Instant information spread model, high randomness and no switch limit.

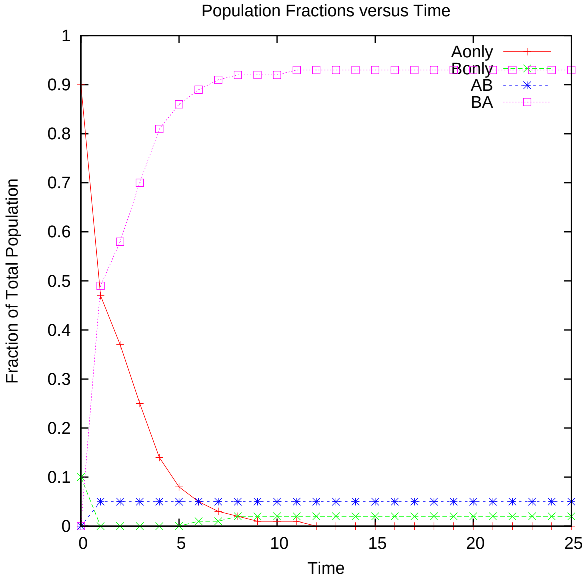

Figure 19: Population distribution when $\frac{\alpha_B}{\alpha_A} = 1.8$ under the Instant information spread model, high randomness and 1 switch limit.

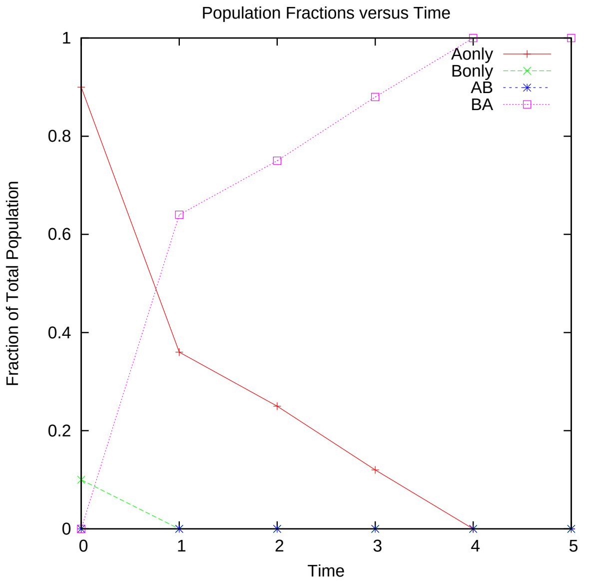

Figure 20: Population distribution when $\frac{\alpha_B}{\alpha_A} = 1.801$ under the ENDOFITER information spread model, no randomness and no switch limit.

Adoption of B is much faster (4-5 epochs) when $\frac{\alpha_B}{\alpha_A}$ is above the threshold even with zero randomness than when below the threshold even with high randomness (15-16 epochs). Above the threshold, limiting the maximum number of switches per user does not affect the system behavior.

The $\frac{\alpha_B}{\alpha_A}$ threshold depends on the initial populations of AONLY and BONLY users. If we decrease the number of AONLY users in the initial population by 1 million, the threshold drops to 1.7. When $\frac{\alpha_B}{\alpha_A}$ is above this lower threshold, all users adopt B under all cases of randomness and in both switching models. However, in all cases, adoption of B is slower than the corresponding case when the threshold was 1.8. Hence, a lower initial population of AONLY users decreases the standalone benefits required to be provided by B for complete adoption. However, the standalone benefits provided by B still needs to be high in order to achieve fast adoption of B.

As $\frac{\alpha_B}{\alpha_A}$ increases, adoption of B is quickened. When the standalone utility of B reaches double that of A, all users switch to BONLY or BA in just 1 iteration even in the ENDOFITER model with zero randomness. In the ENDOFITER model, if there is no

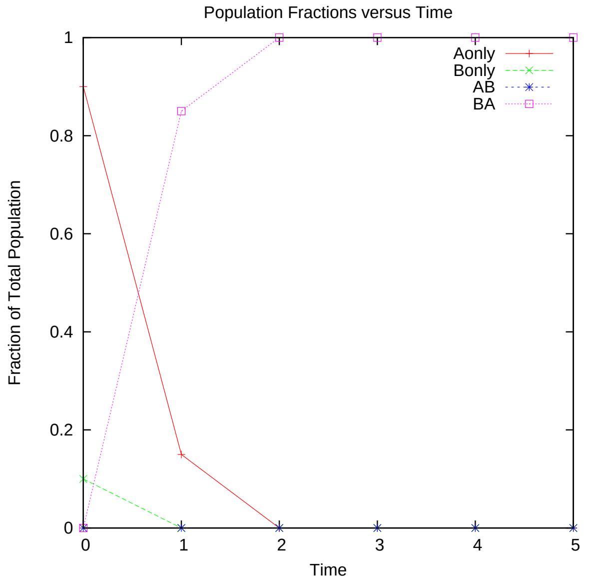

Figure 21: Population distribution when $\frac{\alpha_B}{\alpha_A} = 1.801$ under the INSTANT information spread model, no randomness and no switch limit.

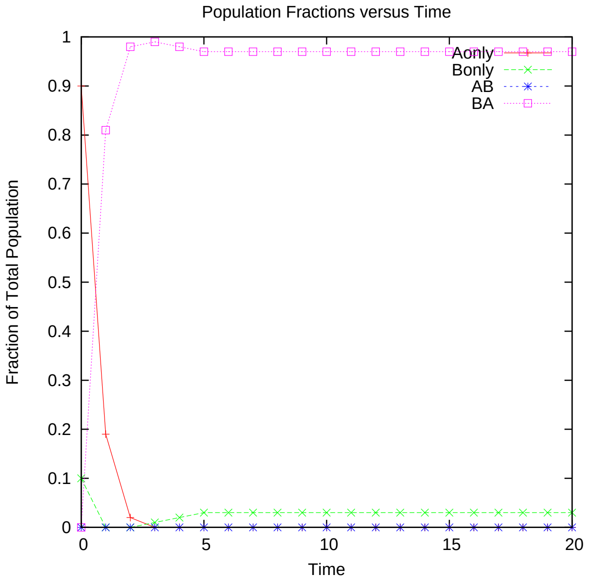

Figure 22: Population distribution when $\frac{\alpha_B}{\alpha_A} = 1.801$ under the INSTANT information spread model, low randomness and no switch limit.

randomness, all users adopt BA converters and never give them up. In the presence of randomness, users first adopt BA converters and then some of them discard the converters, making them BONLY users. Of course, this does not happen if we limit the maximum number of switches per user to 1. Many users forever hold on to their BA converters even in the presence of randomness and no switch limit [2]. This is a result of high switching cost and the lack of degradation on benefits of B caused by the converter ($r_B = 0$). We need to decrease the switching cost associated with discarding BA converters in order to encourage users to become BONLY. This factor must be kept in mind while designing BA converters.

Another way to promote adoption of BONLY is to encourage A users to directly convert to BONLY users without adopting BA converters. This is easier in the INSTANT model – the number of BONLY users never goes to 0 like in other scenarios (Figure 23) However there are still a large number of users unnecessarily hanging on to their BA converters. If we limit the maximum number of switches per user to 2, the fraction of BONLY users at the end of the simulation is higher (Figure 24) – users are prevented from switching back and forth between BONLY and BA.

5.1.1 Take Aways

- There exists a sharp threshold value for $\frac{\alpha_B}{\alpha_A}$, above which complete adoption of B is quick. We should thus strive to increase the standalone utility of B and move $\frac{\alpha_B}{\alpha_A}$ above the threshold. The threshold can be analytically determined using the equations in Section 3.

- A lower initial population of AONLY users decreases the standalone benefits required to be provided by B for complete adoption. However, the standalone benefits provided by B needs to be high in order to achieve fast adoption of B.

- Randomness in switching is beneficial to the adoption of B when $\frac{\alpha_B}{\alpha_A}$ is below the threshold but detrimental when above the threshold.

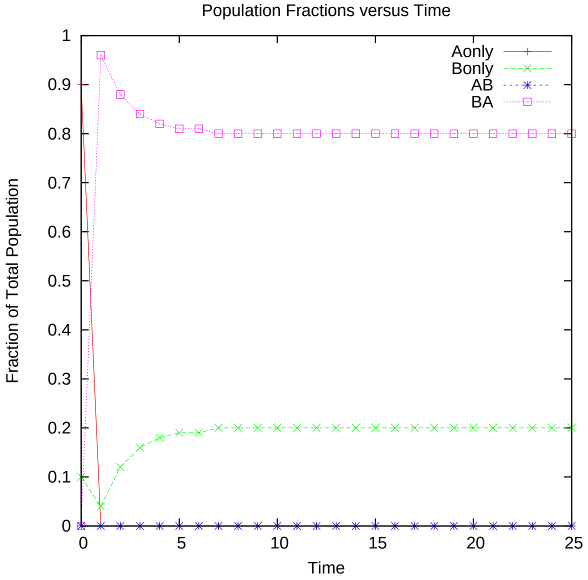

Figure 23: Population distribution when $\frac{\alpha_B}{\alpha_A} = 4.0$ under the INSTANT information spread model, low randomness and no switch limit.

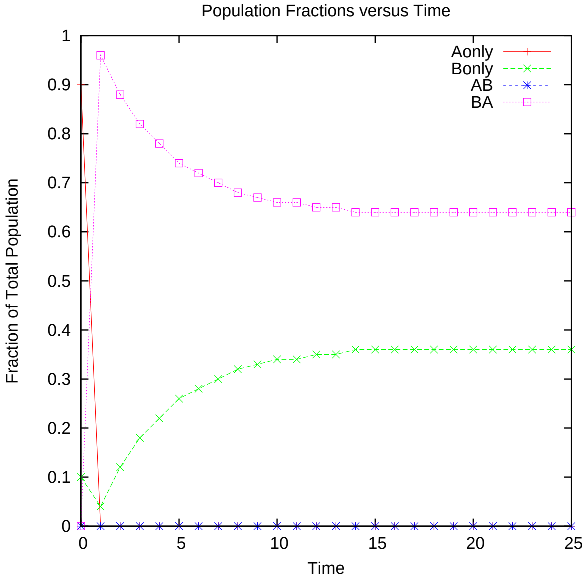

Figure 24: Population distribution when $\frac{\alpha_B}{\alpha_A} = 4.0$ under the INSTANT information spread model, low randomness and 2 switch limit.

---

[2] If there is no switch limit, people keep converting back and forth between BONLY and BA

- Adoption of B is quicker in the INSTANT model than in the ENDOFITER model. Thus, publicizing adoption statistics leads to quicker adoption of the new technology.

- INSTANT information spread aids users to avoid adopting BA converters which they do not have incentive to discard even if the converter become useless. In order to encourage all users to become BONLY, BA converters must be designed to be easily discardable, i.e the switching cost from BA to BONLY is very low. Another way to encourage BA users to discard their converters is to have a positive value for the self-degradation factor r_B. However a positive r_B can be detrimental to the adoption of B during the initial phases.

5.2 Network Effect Parameter β

In this section, we study the impact of network effects by varying β. If there are no network effects, i.e. $\beta = 0$, a large number users (the ones with low switching costs) immediately become BONLY, while others remain AONLY (Figure 25). Note that $\frac{\alpha_B}{\alpha_A} = 1.9$ has been chosen such that switching to B is attractive. In the absence of network effects, the ENDOFITER and INSTANT models behave identically. Limiting the number of switches has no effect either.

When network effects are low, all users slowly switch to BA (Figure 26) if there is no randomness. In the presence of randomness (Figure 27), some of these users become BONLY if more than one switch is allowed. Higher the randomness, more the number of BONLY users. The number of BONLY users is greater when the maximum number of switches is limited to 2 rather than when it is unbounded. This is because back and forth switching between BA and BONLY is prevented by the switch limit. However, limiting the maximum number of switches to 1 in the ENDOFITER model results in a different behavior. Most users become BA, some become AB and very few become BONLY (Figure 28). This is because users get stuck with their initial choices. As expected, adoption of B is faster in the INSTANT model than in the ENDOFITER model.

If the network effects become high (Figure 29), adoption of B does not happen at all – all users adopt A, with or without an AB converter. In the presence of randomness in switching and a maximum switch limit of at least 2, some users convert to AONLY, thus resulting in the system having a mix of AB and AONLY users (Figure 30). The number of AONLY users in the system is higher when the maximum switch limit is 2 rather than ∞ (Figure 31). This is because the maximum switch limit of 2 avoids users from switching back and forth between AONLY and A.

5.2.1 Take Aways

- Small network effects aid in complete adoption of B while zero or very high network effects impede complete adoption.

- Majority of BA users never discard their converters even after all AONLY users have disappeared. When there is randomness in switching, some BA users discard their converters and become BONLY.

- When the system converges to a mix of BONLY and BA users or to a mix of AONLY and AB users, the number of BONLY or AONLY users is more if the maximum number of switches is limited to 2 than when unbounded.

- Complete adoption of B is faster in the INSTANT model than in the ENDOFITER model.

5.3 Converter Efficiency q

Converters typically offer only a fraction of the benefits of the other technology. In this section, we analyze how converter efficiencies impact the adoption of B. In the absence of converters, all users switch to AONLY in order to reap the network benefits associated with the large initial population of AONLY users (Figure 32)[3]. Adoption of B has no chance in

---

[3] Note that we are using the default value of $\frac{\alpha_B}{\alpha_A}$ = 1.8

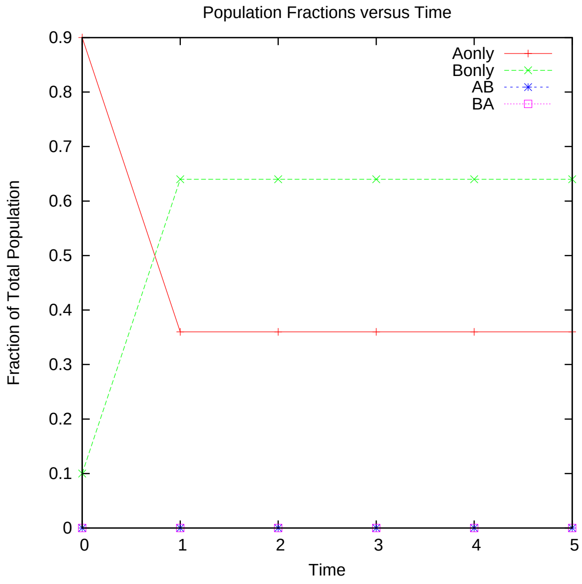

Figure 25: Population distribution when $\beta = 0$ and $\frac{\alpha_B}{\alpha_A} = 1.9$ in the EndOfIter model, no randomness and no switch limit.

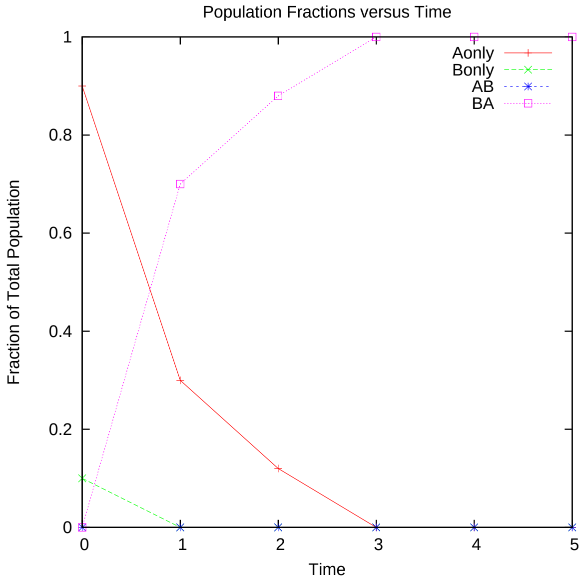

Figure 26: Population distribution when $\beta = 0.001$ and $\frac{\alpha_B}{\alpha_A} = 1.9$ in the EndOfIter model, no randomness and no switch limit.

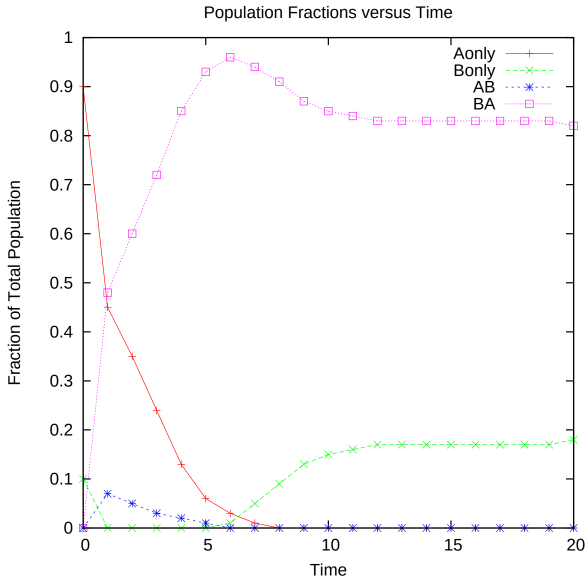

Figure 27: Population distribution when $\beta = 0.001$ and $\frac{\alpha_B}{\alpha_A} = 1.9$ in the EndOfIter model, high randomness and no switch limit.

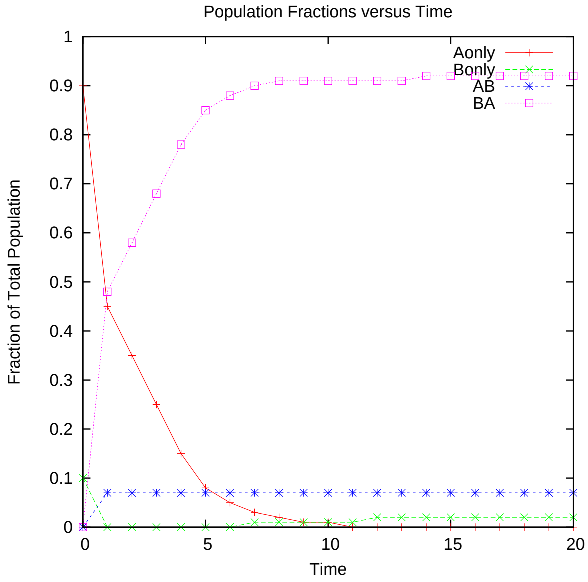

Figure 28: Population distribution when $\beta = 0.001$ and $\frac{\alpha_B}{\alpha_A} = 1.9$ in the EndOfIter model, high randomness and 1 switch limit.

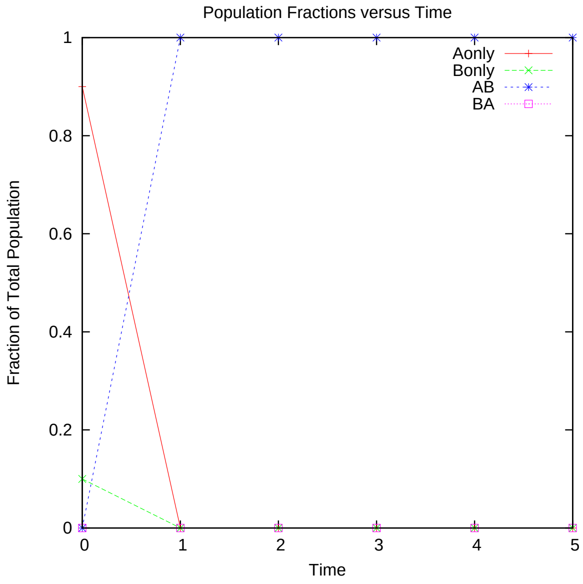

Figure 29: Population distribution when $\beta = 0.01$ and $\frac{\alpha_B}{\alpha_A} = 1.9$ in the ENDOFITER model, zero randomness and no switch limit.

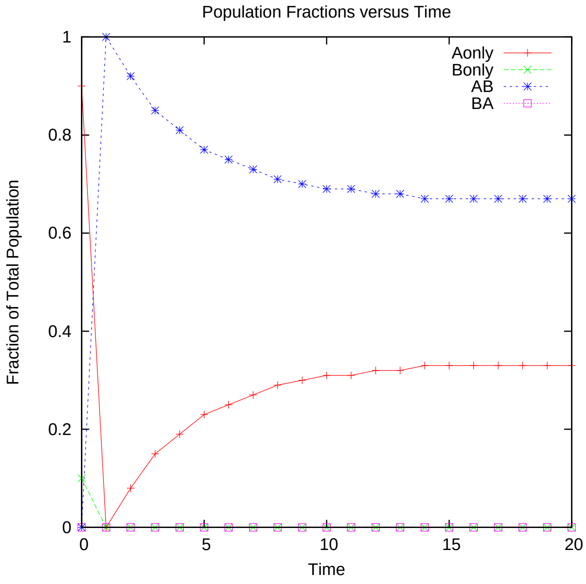

Figure 30: Population distribution when $\beta = 0.01$ and $\frac{\alpha_B}{\alpha_A} = 1.9$ in the ENDOFITER model, high randomness and 2 switch limit.

this scenario unless we increase factors like the standalone utility of B, as analyzed earlier or we deploy converters.

When both AB and BA converters have low efficiency (10%), users switch to AB or remain AONLY (Figure 33). The little extra network effects enabled by the inefficient BA converter is not enough to overcome the huge network effects obtained by being an AONLY or AB user in the large population of AONLY users.

All users adopt B when both AB and BA converters are 90% efficient (Figure 34). In the absence of randomness in switching, all users adopt BA converters. If there is randomness in switching, the system converges to a mix of BONLY and BA users (Figure 35). Higher the randomness, higher is the number of BONLY users. Randomness also delays the complete adoption of B. This is because of the random back and forth switching between BONLY and BA. Adoption of B is faster in the INSTANT model than in the ENDOFITER model (Figure 36) – 6 iterations versus 8 iterations.

If the BA converter is 100% efficient and the AB converter is only 10% efficient, in the ENDOFITER model, all users switch to BA immediately (Figure 37). However, in the INSTANT model, the number of AONLY users never goes to 0 and complete adoption of B never takes place (Figure 38). As the BA converter is 2-way, even AONLY users can benefit from BA converters adopted by previously BONLY users. So A users are less inclined to switch as they immediately come to know about the extra network effects from the BONLY to BA converts. Under high randomness, complete adoption of B takes longer in the INSTANT model than in the ENDOFITER model. Hence, depending on the BA converter efficiency, INSTANT information spread is sometimes beneficial to the adoption of B while it is detrimental at some other times.

When the AB converter is 10% efficient and the BA converter is 92.5%, all users adopt B immediately in the ENDOFITER model. Complete adoption of B happens faster when the BA converter efficiency is 98%. However, if the BA converter efficiency is 94.5%, complete adoption of B does not

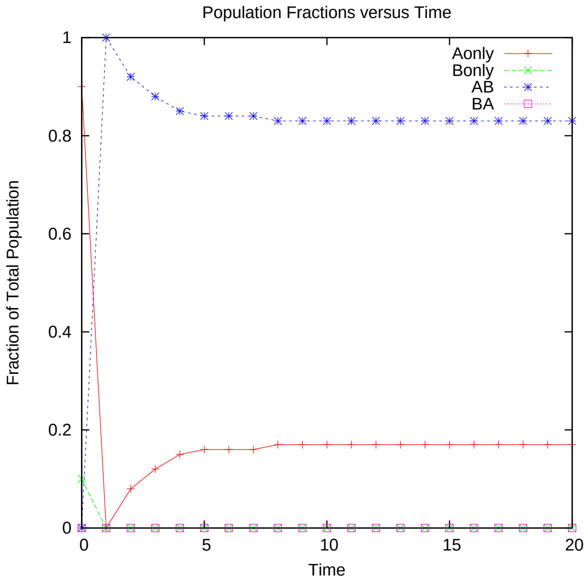

Figure 31: Population distribution when $\beta = 0.01$ and $\frac{\alpha_B}{\alpha_A} = 1.9$ in the EndOfIter model, high randomness and no switch limit.

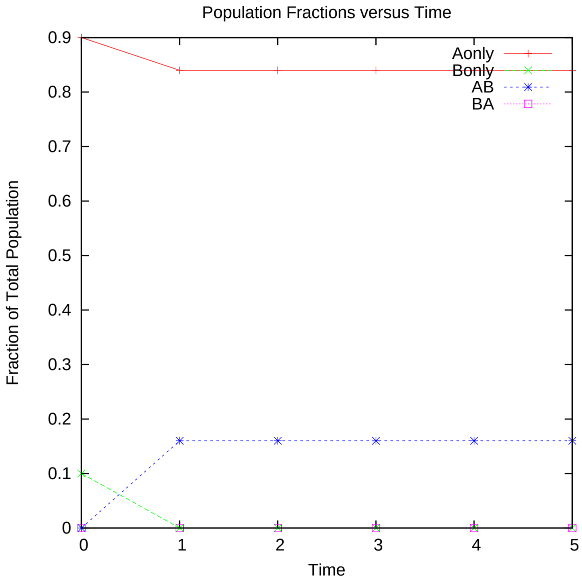

Figure 33: Population distribution with 10% efficient converters in the EndOfIter model, 0 randomness and no switch limit.

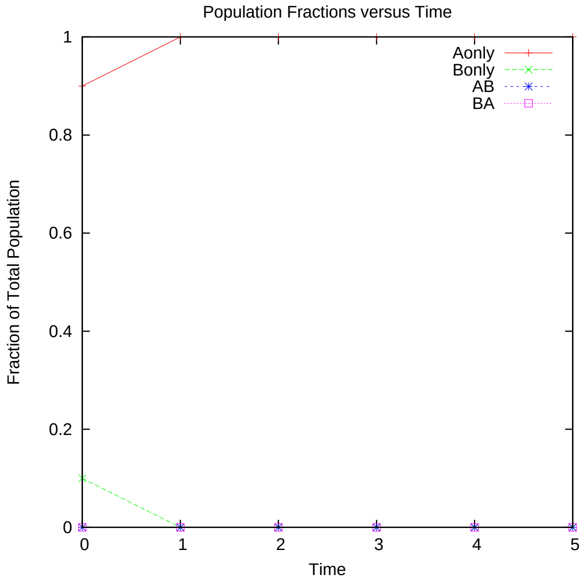

Figure 32: Population distribution when no converters are present in the EndOfIter model, 0 randomness and no switch limit.

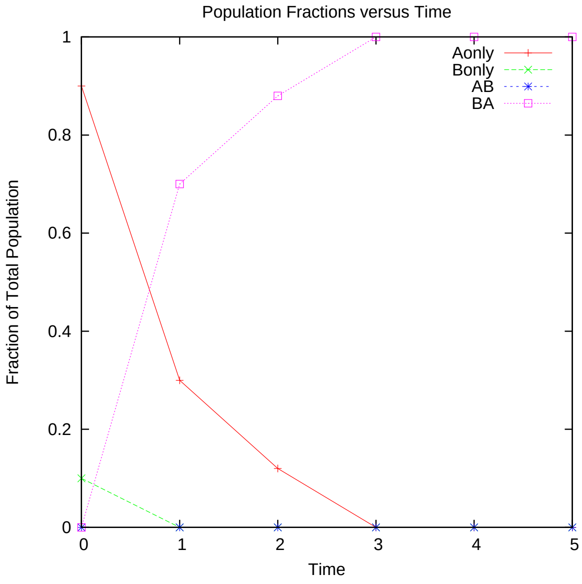

Figure 34: Population distribution with 90% efficient converters in the EndOfIter model, 0 randomness and no switch limit.

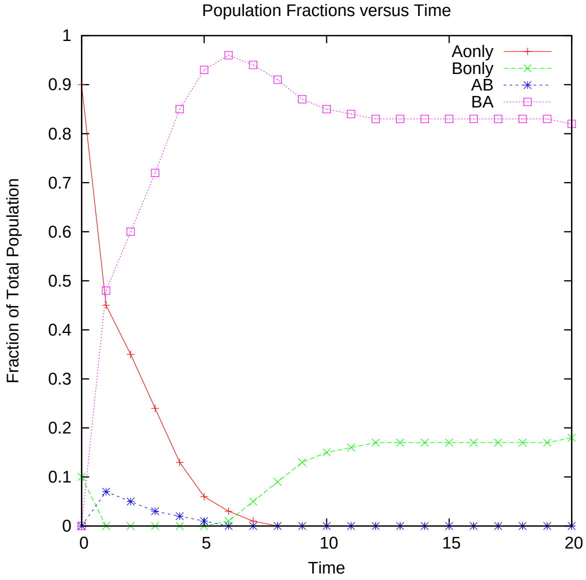

Figure 35: Population distribution with 90% efficient converters in the ENDOFITER model, high randomness and no switch limit.

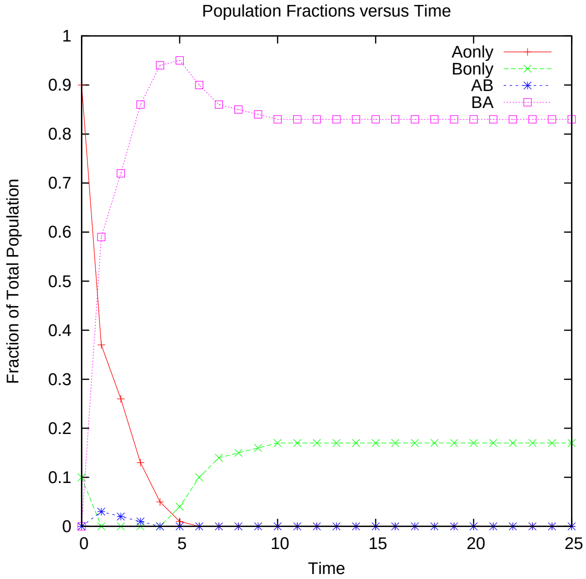

Figure 36: Population distribution with 90% efficient converters in the INSTANT model, high randomness and no switch limit.

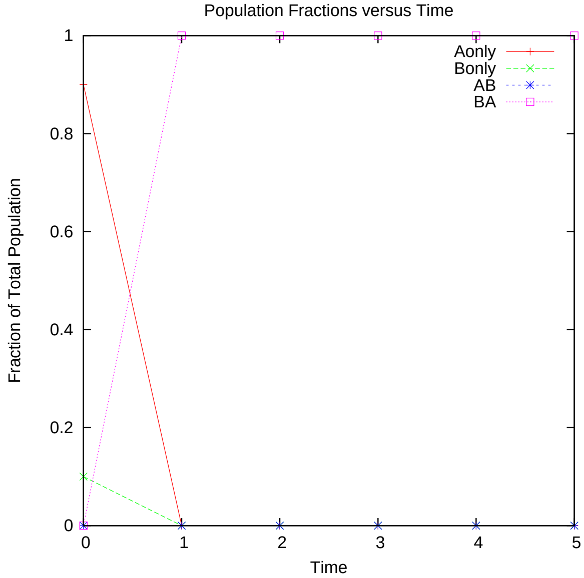

Figure 37: Population distribution with 100% efficient BA converter and 10% efficient AB converter in the ENDOFITER model, 0 randomness and no switch limit.

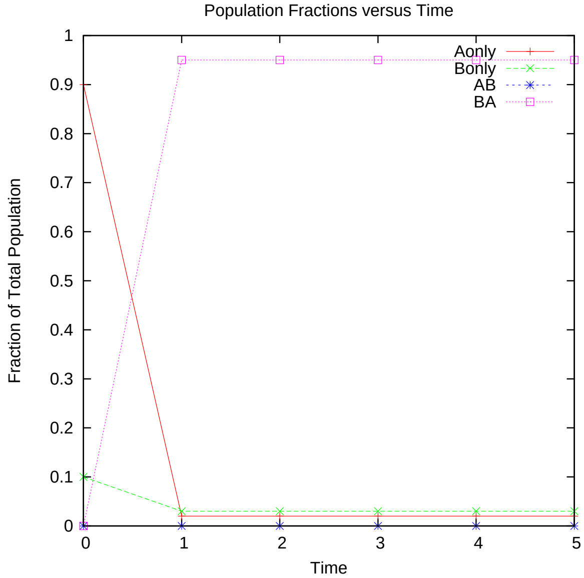

Figure 38: Population distribution with 100% efficient BA converter and 10% efficient AB converter in the INSTANT model, 0 randomness and no switch limit.

take place at all. Figure 39 shows how the fraction of B users varies with time for the three values of converter efficiency under consideration. When the converter efficiency is 94.5%, the penetration of B gets stuck at around 95%. This behavior occurs even in the presence of maximum switch limits, but does not happen when there is randomness or in the INSTANT model. Further analysis found that adoption of B progresses via AONLY users switching to BA only if the difference in utilities between BA and AONLY increases at the end of each iteration. For this to hold true at the end of an iteration, the following condition (derived from equations in Section 3) must hold true:

$$2N_{A->Bswitches} \geq N_{Bonly}(1/q_A - 1) \qquad (9)$$

At high values of BA efficiency (i.e. lower q_A), a larger number of A to B switches are required to prevent the adoption of B from stalling. Although, higher converter efficiency leads to higher number of switches, other model parameters influence the number of switches and can make it fail to satisfy Equation 9. This happens when the BA converter is 94.5% under model parameter values used in the simulation. Hence, converters must be carefully engineered such that their efficiencies are neither too low nor too high to promote quick and complete B adoption.

5.3.1 Take Aways

- When the BA converter efficiency is very high, adoption of B is faster in the ENDOFITER model than in the INSTANT model.
- Higher BA converter efficiency does not always lead to quicker adoption of B.There exists values e_1, e_2 and e_3 for BA converter efficiencies ($e_1 < e_2 < e_3$) such that complete adoption of B happens under e_1 and e_3, but stalls under e_2. Thus, converters must be engineered to have the right value of efficiency.

5.4 Performance Hit due to Converter r

A converter enables a user to communicate with users belonging to a different technology and reap the network benefits. But in many cases, running a converter detrimentally affects the benefits provided by the user's current technology choice. For example, running an IPv4 to IPv6 converter may expose the user to more attacks due to outdated firewall rules. Running the OCALA proxy on a user's desktop may decrease the desktop's computation and network performance due to packet interception by the OCALA proxy. In this section, we study the impact of the performance hit caused by converters on the adoption of the new technology by varying the parameters r_A and r_B.

When neither the AB nor the BA converter causes any performance degradation, i.e. $r_A = r_B = 0$, all users adopt BA converters (Figure 40). As seen in results earlier, adoption of B is slower in the presence of randomness while the INSTANT model quickens it. The system converges to a mix of BONLY and BA users when randomness is present or in the INSTANT model. Higher the randomness, more the number of BONLY users in the mix.

When both AB and BA converters cause a 3.25% hit in the benefits provided by A and B respectively, complete adoption of B still occurs (Figure 41), albeit 16 times slower than in the previous case. As the number of BA users increases, the individual utility associated with a BONLY user keeps increasing and goes higher than the individual utility associated with a BA user (Figure 42). At this point, users with small switching costs discard their BA converters and become BONLY users. Although the number of BONLY users when the system converges is greater than the number of BONLY users in the initial population, majority of the users are still BA. The performance hit caused by the BA converter is not large enough to prompt the majority of the population to discard their BA converters which are now useless in the absence of AONLY users. In the presence of high randomness, there are more BONLY users than BA users when the system converges (Figure 43). The INSTANT model greatly speeds up the adoption of B - complete adoption of B in the INSTANT model takes approximately one-fifth the time taken in the ENDOFITER model (Figure 44). Adoption of B takes a very long time irrespective of any limit on the maxi-

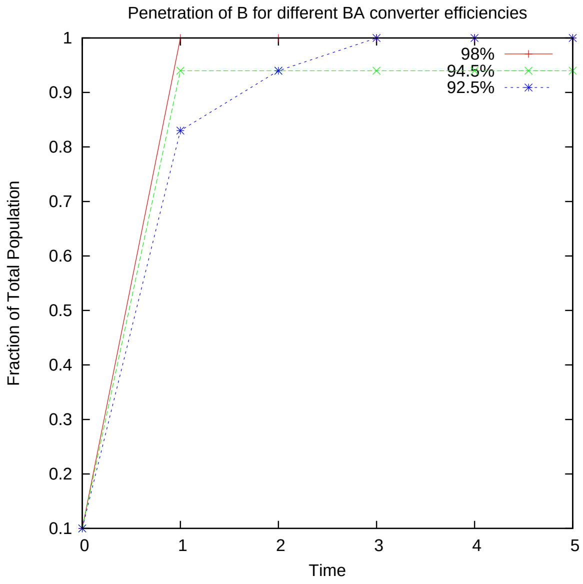

Figure 39: Penetration of B over time for different values of BA converter efficiencies.

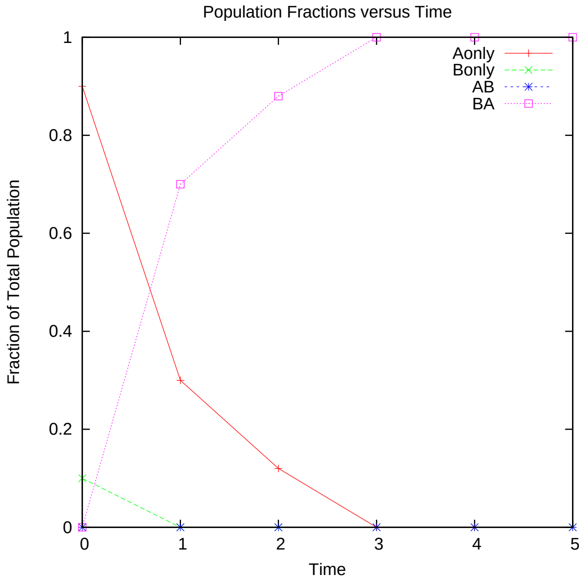

Figure 40: Population distribution when $r_A = r_B = 0$ in the EndOfIter model, 0 randomness and no switch limit.

mum switches. Even if we limit the maximum number of switches to 1 we see some Bonly users when the system converges. The high performance hit of the BA converter made these users directly jump to Bonly unlike in previous scenarios when users first convert to BA and get stuck there.

If we increase the performance hit of the converters to 5%, complete adoption of B stalls in the zero randomness scenario (Figure 45). Under high randomness, adoption of B still proceeds to completion, resulting in a mix of BA and Bonly users, with Bonly users in the majority.

In all cases considered so far, both AB and BA converters cause an equal performance hit. If only the AB converter causes a performance hit, the results remain the same. If only the BA converter causes a hit, adoption of B stalls as most users remain with A and adopt the superior AB converters.

5.4.1 Take Aways

- Users are more encouraged to discard their BA converters if the performance hit on the benefits provided by B is high. However if the performance hit is too high, adoption of B stalls.

- As the performance hit caused by the BA converter increases, the time for complete B adoption increases. Adoption of B is faster in the Instant model than in the EndOfIter model by a factor of 16.

5.5 Switching Costs

We study the impact of the switching cost on the adoption of B by varying the range from which the switching costs of individual users are picked. When the switching cost is negligible, all users adopt B immediately (Figure 46). If there is randomness in switching and no limits on the maximum number of switches per user, equal number of Bonly and BA users are present during convergence Randomness in switching leads to users continuously switching back and forth between Bonly and BA. Randomness in switching is thus highly detrimental to convert all users to Bonly, especially when the switch-

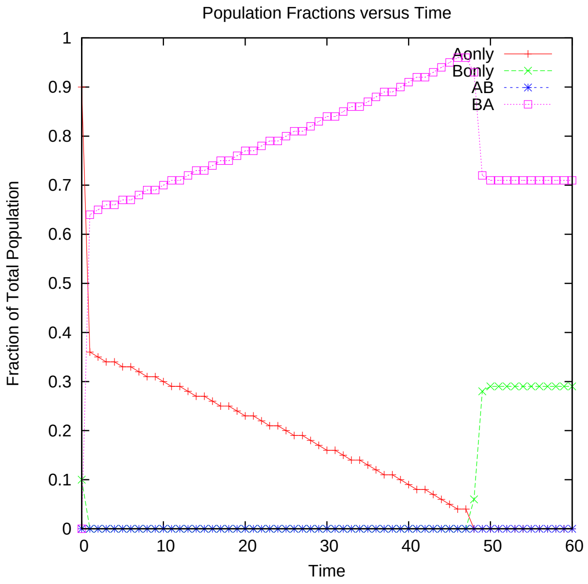

Figure 41: Population distribution when $r_A = r_B = 0.0325$ in the ENDOFITER model, 0 randomness and no switch limit.

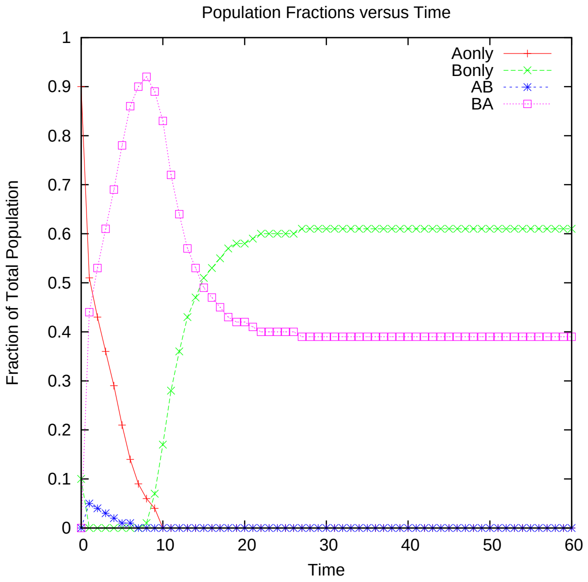

Figure 43: Population distribution when $r_A = r_B = 0.0325$ in the ENDOFITER model, high randomness and no switch limit.

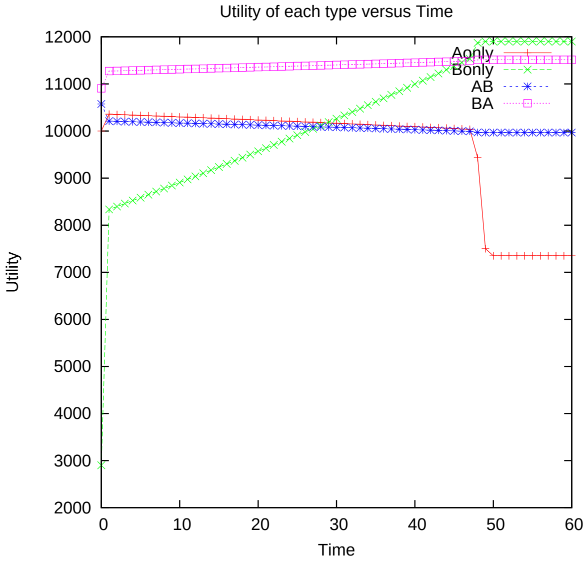

Figure 42: Individual utilities of the various technology types when $r_A = r_B = 0.0325$ in the ENDOFITER model, no randomness and no switch limit.

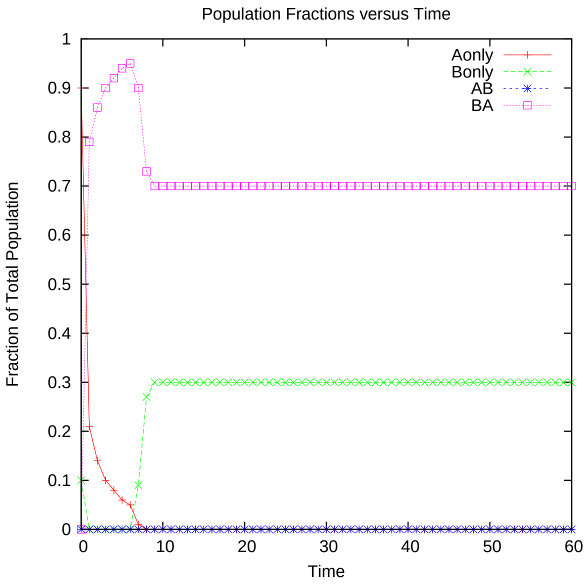

Figure 44: Population distribution when $r_A = r_B = 0.0325$ in the INSTANT model, 0 randomness and no switch limit.

ing costs are low. (Figure 47). Under the ENDOFITER model, all users convert to BA before becoming BONLY. In the INSTANT model, the number of BONLY users never goes to 0 – some users directly convert from AONLY to BONLY. If the maximum number of switches per user is limited to 2, all users convert to BONLY (Figure 48).

As the switching costs increase, adoption of B slows down as expected (Figure 49). Adoption of B is faster in the INSTANT model than in the ENDOFITER model. Randomness in switching delays the adoption of B. If the switching cost becomes very high, adoption of B gets stalled (Figure 50).

5.5.1 Take Aways

- Increased switching costs slow down the adoption of B.
- All users convert to BONLY if the maximum number of switches limited to 2. Under very low switching costs, if there is no limit on the number of switches, the system converges to an equal mix of BONLY and BA users.
- As in prior sections, adoption of B is faster in the INSTANT model than in the ENDOFITER model. Randomness in switching slows down the adoption of B.

5.6 External Benefits of the Converter, t

A BA converter often provides a fraction of the standalone benefits of A in addition to the network benefits of being able to communicate with AONLY users. The same applies for AB converters. In this section, we study the impact of this feature of a converter on the adoption of B, by varying the parameters t_A and t_B.

When both AB and BA converters provide a small fraction ($t_A = t_B = 0.01$) of the standalone benefits of the other technology, all users switch to BA (Figure 51). Like in previous results, randomness slows down the complete adoption of B and adoption is faster in the INSTANT model than in the ENDOFITER model. In the ENDOFITER model, all users first become BA and then some of them discard their con-

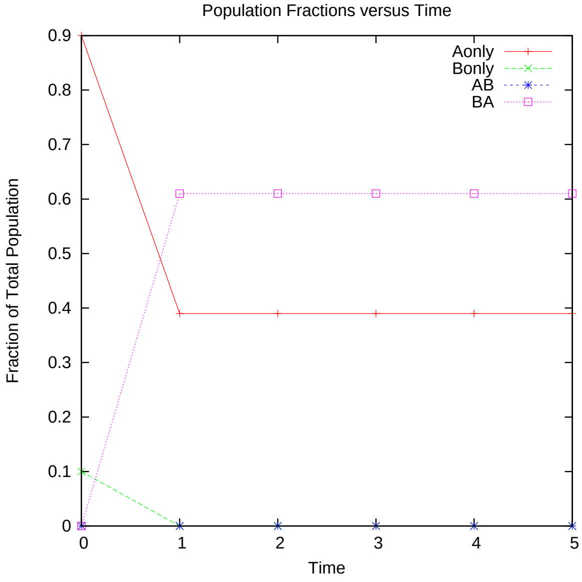

Figure 45: Population distribution when $r_A = r_B = 0.05$ in the ENDOFITER model, 0 randomness and no switch limit.

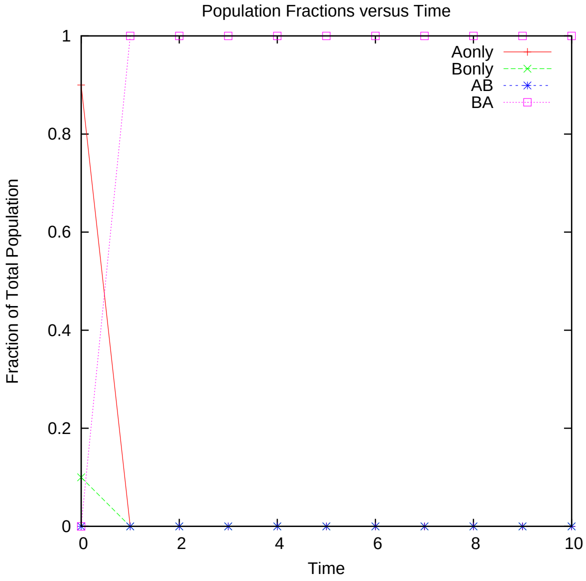

Figure 46: Population distribution when the switching cost range is 0-1 in the ENDOFITER model, 0 randomness and no switch limit.

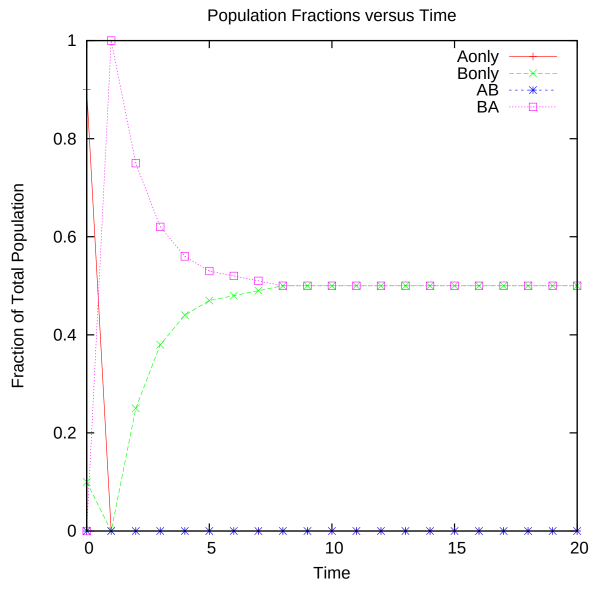

Figure 47: Population distribution when the switching cost range is 0-1 in the ENDOFITER model, high randomness and no switch limit.

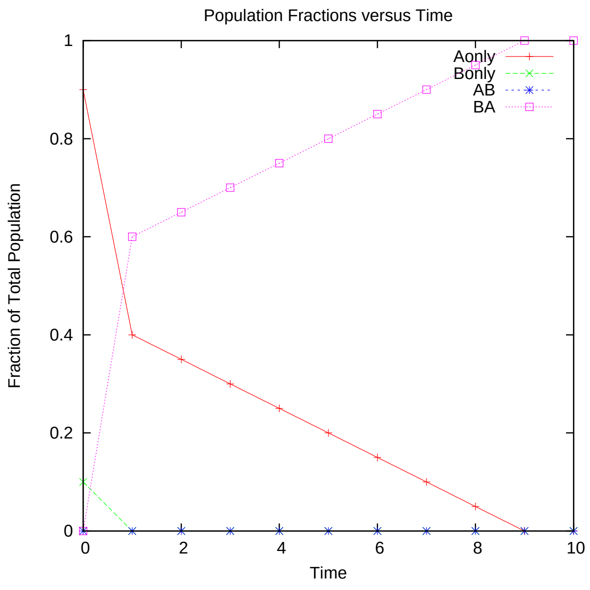

Figure 49: Population distribution when the switching cost range is 0-1800 in the ENDOFITER model, 0 randomness and 0 switch limit.

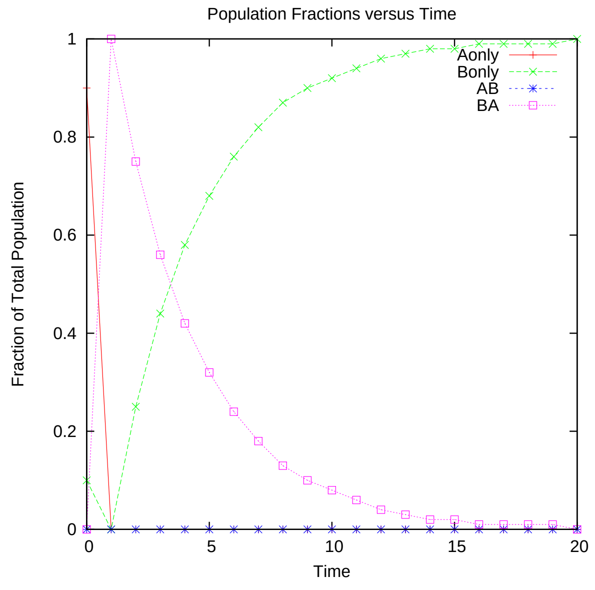

Figure 48: Population distribution when the switching cost range is 0-1 in the ENDOFITER model, high randomness and 2 switch limit.

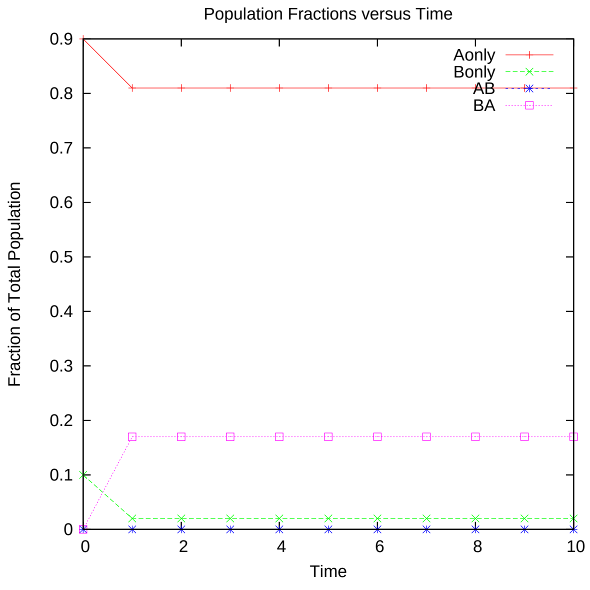

Figure 50: Population distribution when the switching cost range is 0-10000 in the ENDOFITER model, 0 randomness and 0 switch limit.

verters to become BONLY. In the INSTANT model, some users directly convert to BONLY; the number of BONLY users in the system never goes to 0. If the number of switches is limited to 1, a very small fraction of users get stuck with AB converters, AB converters are now more attractive.

In the ENDOFITER model, complete adoption of B takes place as long as t_A and t_B are below 11.111111%. This threshold was derived by analyzing the equations in Section 3. Above this threshold, the individual utility associated with AB is greater than that associated with BA and complete adoption of B does not happen (Figure 52). All users switch to AONLY or AB. The number of AB users in the INSTANT model is less than that in ENDOFITER model. Users directly convert to AONLY and avoid converting to AB as the information about the diminishing number of BONLY users reaches their ears more quickly in the INSTANT model. Under very high randomness in switching, complete adoption of B takes place even above the threshold value.

In the previous case, both AB and BA converters offered the same fraction of the standalone benefits of B and A respectively. If only the BA converter offers the extra benefits, adoption of B is easier as BA converters are more attractive than AB converters. If only AB converters offer the extra benefits, complete adoption of B still happens if the percentage of the standalone benefits of B offered by the AB converter is less than 5.26%. This threshold was derived from the equations in Section 3 and verified by simulation (Figures 53 and 54).

5.6.1 Take Aways

- Complete adoption of B happens as long as the fraction of standalone benefits of the competing technologies offered by the AB and BA converters is less than a threshold.

- If only the AB converter offers a fraction of the standalone benefits of B, complete adoption of B still happens if this fraction is lower than a threshold. The value of this threshold is lower than the threshold when both AB and BA converters offer an equal fraction of the standalone

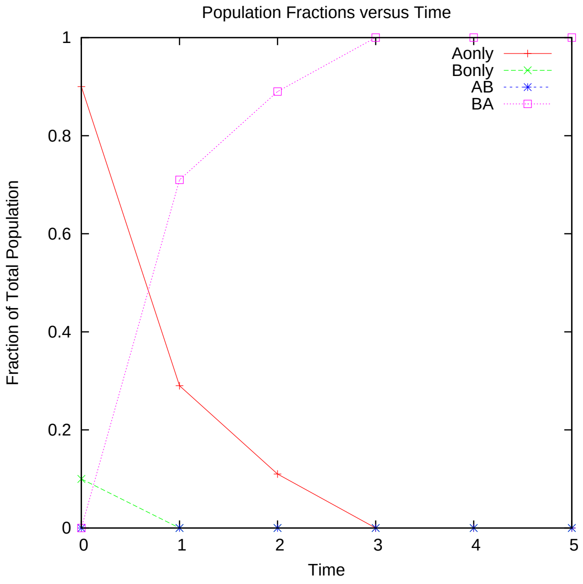

Figure 51: Population distribution when $t_A = t_B = 0.01$ the ENDOFITER model, 0 randomness and no switch limit.

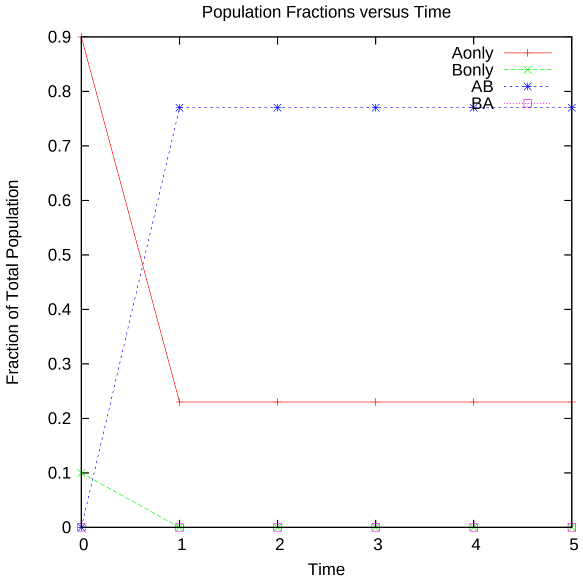

Figure 52: Population distribution when $t_A = t_B = 0.11111112$ in the ENDOFITER model, 0 randomness and no switch limit.

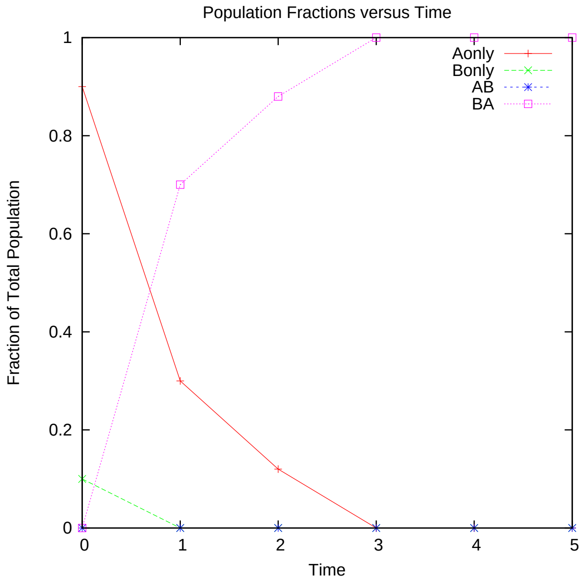

Figure 53: Population distribution when $t_A = t_B = 0.0526$ in the ENDOFITER model, 0 randomness and no switch limit.

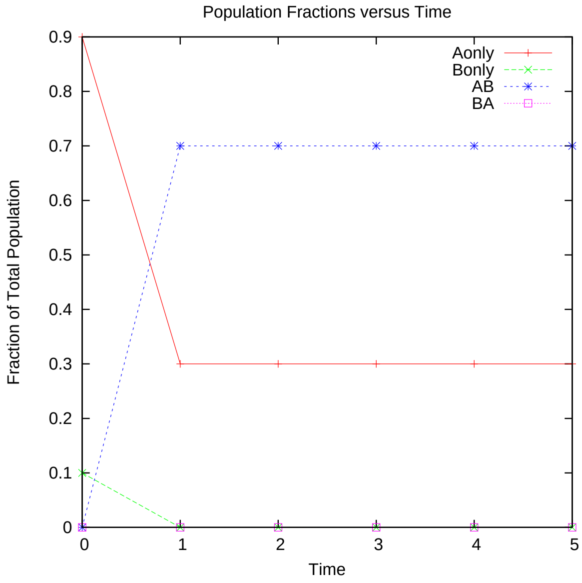

Figure 54: Population distribution when $t_A = t_B = 0.0527$ in the ENDOFITER model, 0 randomness and no switch limit.

benefits of the respective competing technologies.

5.7 INSTANT **versus** ENDOFITER

In this section, we summarize the impacts of the ENDOFITER and INSTANT information spread models seen so far. Complete adoption of B is faster under the INSTANT information spread model than under the ENDOFITER model. However, if the BA converter is close to 100% efficient, the reverse is true. This is because converters are two-way and INSTANT information spread quickly informs A users about the increased number of BA users they can 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, information about B adoption should be publicized and quickly dispersed to all users.

6 Limitations

This report is only a first step in modeling and analytically studying the adoption of new network architectures. Our current model of new network architecture adoption is very basic and simple. Many missing aspects like interactions between ASes, infrastructure vendors, application vendors and organizations, different switching costs and utility functions for different users, economies of scale and learning as penetration increases and impact of geo-political boundaries and influences are part of our ongoing research agenda. The parameter values used in our simulation studies do not directly map onto real-world numbers. We attempt only to draw general conclusions about the relative importance of the model parameters and how they relate to the real world.

7 Conclusion

Studying the adoption of new network architectures through mathematical analysis and simulations is a fruitful research area. It corroborates intuitive results 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 non-obvious results like higher efficiency converters do not always aid the adoption of the new network architecture. This report increases our understanding of the factors that influence adoption and gives insights into how we can improve network architecture design and engineering to achieve quick adoption of a new network architecture like IPv6.

References

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

[2] IPv6 economic impact assessment - final report. National Institue of Standards and Technology, U.S. Dept. of Commerce, Oct 2005.

[3] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Resilient Overlay Networks. In *SOSP 2001*.

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

[5] 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.

[6] Farrell, Joseph and Saloner, Garth. Converters, compatibility, and the control of interfaces. *The Journal of Industrial Economics*, 40(1):9–35, mar 1992.

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

[8] 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*.

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