On Economic Viability and Choices of Networked Systems & Architectures

Soumya Sen
ESE, U.Pennsylvania

1. INTRODUCTION

Networked systems and architectures have a ubiquitous presence in today's world. The Internet, electrical power grids, facilities management networks etc., are just a few examples of such systems that permeate our daily lives. The potential for both economic and technological growth offered by these networked systems have attracted businesses to invest in them, spurring further research. This kind of networking research is a fast emerging multi-disciplinary field that draws from diverse disciplines, such as computer science, economics, and sociology - a convergence driven by the fact that success or failure of a technology depends not only on its scientific merits but also on many complex socio-economic factors. Over the years there have been many network technologies that met their technical specifications, and yet failed to become economically viable. For example, the failure of wide-spread deployment of QoS architectures and solutions in today's Internet is most commonly attributed to economic reasons, like the lack of user demand, high operational costs (compared to over-provisioning), and inter-provider settlement issues. Another such problem is also being encountered in the migration from IPv4 to IPv6. The lesson is that in order to develop technologies that are eventually successful, researchers need to make design choices which account for different economic aspects, such as supply and demand side uncertainties, consumer preferences, weight of incumbency, competition strategies, costs etc. However, in most cases it is often unclear which of these can have a really significant impact. Therefore, developing models that include economic as well as technical factors is of major importance. Additionally, these models can also be very useful to business strategists, policy makers, and social planners alike for analyzing questions regarding corporate and social welfare.

The goal of this research is to identify how various economic factors influence design choices and tradeoffs in networked systems and architectures, their deployment and adoption, and thus build a framework for reasoning about choices and decision-making. This framework needs to incorporate different models which address specific issues and questions that network technology providers commonly face, such as, how can a new network technology compete better against an incumbent, or how will a new technical capability (*e.g.,* virtualization) affect network profitability, etc.?

These questions are clearly far-reaching and multi-faceted, and in the context of this dissertation, we concentrate on three basic issues that are important from a network provider's perspective. First is the issue of understanding how a provider's decisions and actions impact the process of technology adoption, and in particular, the migration from an incumbent to an entrant technology. A second related issue is that of choosing what type of network architecture should the entrant technology be deployed on, and the third issue is to understand how the range of functionalities included as a part of the network architecture affects its profitability and potential for future innovations. These are the three broad topics that we consider in our research, and they are briefly discussed next.

As new network technologies and services become available, the question that their providers ask is whether these technologies will be successful. The main challenge faced by most new network technologies is the weight of incumbency, *i.e.,* market penetration of an existing technology against which the entrant needs to compete. The presence of a large installed base can give an incumbent an edge even if an entrant is technically superior. This is because the benefit that a user derives from using a network technology grows proportionally with the network size, *i.e.,* number of other users reachable on it; a feature also known as positive *network externality* effect [7]. Network providers have traditionally relied on converters (a.k.a. gateways) to solve this problem. Converters help entrants overcome the influence of the incumbent's installed base by enabling cross-technology inter-operability. However, converters can also introduce performance degradations and functionality limitations, and may even end up helping the incumbent by mitigating the impact of its users migrating over to the new technology. Thus, it is often unclear to the provider as to when and how converters can facilitate network technology migration. Additionally, there are also factors like pricing, switching costs, user preferences etc., which impact the potential for a successful migration. Having access to economic models that incorporate these factors will allow providers to quantify these impacts and take actions to improve the chances of their technology's success. To this end, we develop a model for adoption of competing network technologies by individual users, and use it to analyze the influence of various economic factors, and in particular, the role of converters in the migration from an incumbent to a new technology. Subsection 2.1 on '*Network Technology Migration*' provides further details regarding this model and the insights gained from it.

While the previous topic focuses on how network technologies compete for market penetration, a related issue for the provider is to select the network architecture on which they are to be offered. Providers of these new technologies or services have to decide whether multiple services should be deployed on a common network infrastructure through resource-sharing, or should each service be offered on a dedicated network of its own. While the option of sharing resources helps realize economies of scope in set-up costs etc., combining heterogeneous services on a single network is not necessarily the right option because of the increased complexity in managing services with disparate requirements. Although analyzing these economic trade-offs is crucial, it is not enough for choosing the right architecture. This is because an architecture's profitability is also influenced by the demand uncertainty of new services. The presence of uncertainty can not only influence resource allocation decisions but also the very choice of architecture. Moreover, recent technological trends (*e.g.,* virtualization) which give providers the ability to dynamically adjust resources in response to changes in demand further adds to the decision's complexity. Therefore, it is important to analyze how these factors can affect the choice between the shared and dedicated architecture options. We propose to develop an economic model that captures all these different aspects of the provider's decision problem. Subsection 2.2 on '*Shared versus Dedicated Network Architecture*' will present the model along with our findings, and discuss how they translate into guidelines for choosing the right architecture.

The third topic that we address in this research deals with how much of functionalities or capabilities do providers need to incorporate in their network architecture. This issue is particularly important in the context of open network architectures where *infrastructure providers* create the network architecture with built-in functionalities, and *service providers* (*e.g.,* Yahoo, Google) use these functionalities to develop and deploy new services on this architecture for end-users to access. There is a natural tradeoff between creating a functionality-rich versus minimalist network architecture. A network with very little functionality will potentially have a few services running on it because the cost of developing these additional functionalities will need to be borne by the service providers. The limited number of services will also make the network less attractive to end-users. Therefore, the fees that the infrastructure provider can charge to the two sides (service providers and end-users) will have to be kept relatively low, thus adversely affecting the network's profitability. On the other hand, if the infrastructure provider were to invest very heavily in creating a functionality-rich architecture, then to compensate for these expenses it will have to charge high fees to the service providers and/or end-users, resulting in fewer subscriptions and reduced profits. Additionally, these high fees charged to the service providers will also discourage them from innovating better functionalities of their own, and instead use the built-in ones, irrespective of their quality. Thus, it seems plausible that the network's richness in functionality may end up stifling service innovation. Hence, it becomes important to analyze these trade-offs within an analytical framework to better understand if and when should network architectures with greater functionalities be created, or conversely, whether there is an 'optimal' level of functionalities that networks should have. This is a topic of our ongoing work and we discuss the initial formulation this problem as a two-sided market model in subsection 2.3 on '*Minimalist versus Functionality-rich Network Architecture*'.

Before proceeding with each individual topic, we wish to note that there are two 'big picture' questions that are fairly common across all these topics: (i) how should we account for technical factors (*e.g.,* technology quality, efficiency, architectural flexibility) and economic factors (*e.g.,* user demand, costs, market mechanisms) while making decisions?; (ii) how do these decisions affect outcomes like network profitability and economic dynamics, such as competition, demand, pricing and costs? The first question requires identifying the various features relevant to the topic and then developing economic models that incorporate them. These aspects will be discussed under the *Problem Description* and *Model Description* subsections of each topic. The second question requires that these models be analyzed to extract general economic conclusions and use them to offer guidelines for decision making. These will be presented in the subsection on *Key Results* for each topic.

2. RESEARCH TOPICS

2.1 Network Technology Migration

2.1.1 Problem Description

Networks, like other technologies, constantly improve over time as newer and better solutions become available. These entrant technologies compete against the incumbents they seek to eventually replace. The Internet itself is an example of such a technology that competed against other alternative packet data networks to finally displace the traditional phone network as the de facto communication infrastructure. But such successful migration from an incumbent network technology to an entrant depends not only on the latter's technical superiority but also on economic factors, and the ability to win over the incumbent's installed base.

The traditional networking approach to facilitate technology migration has been the introduction of converters or gateways. Converters can help a technology to increase its network externality benefits by allowing its users to connect with users of the another technology. However, developing, deploying, and operating converters come at a cost, one that often grows as a function of the converter's quality. Further, converters can play a directionally ambiguous role. On one hand, converters can help the entrant overcome the advantage of the incumbent's large installed base by allowing connectivity to it. But on the other hand, they also help the incumbent technology by mitigating the impact of its users migrating to the newer technology.

To help network technology providers better understand how the deployment of converters and other economic factors influence network technology migration, we develop an economic model which incorporates these factors (*e.g.,* quality, externality, price). We consider the utility derived from network technologies by individual heterogeneous users, and use it to build an aggregate model for technology adoption that is consistent with individual rational decision-making.

Next, we briefly discuss how the adoption model we propose in this work compares to the existing literature. Modeling the adoption of new products and technologies has a long tradition in marketing. Fourt and Woodlock [13] proposed a product diffusion model in which a fixed fraction of consumers who have not yet bought the product did so at every period; this is also known as constant hazard rate model. Bass [1] extended it to incorporate word-of-mouth communication between current adopters and potential buyers. These earlier single technology adoption models were later extended to study the joint diffusion of successive generations of technologies by Norton & Bass [11]. However, the focus of all these works is on the aggregate adoption dynamics as opposed to modeling the individual user's decision-making process. As a result, these models fail to develop an understanding of how the consumer decision process affects adoption dynamics and how various economic factors impact adoption decisions. Only a few models have focused on individual-level adoption [2], which provide much greater insights into the mechanism through which rational individual decisions result in aggregate system dynamics. Due the complexity of such models, their use has been limited to the case of a single technology. We have extended these models to a two technology settings -an essential step towards making them suitable for studying migration from an incumbent to an entrant network technology and for investigating the role that converters play in that process. Other prior works that have considered converters in the adoption of incompatible technologies include Farrell & Saloner [4, 12], Choi [3], Joseph *et al.* [6]. The key findings of these works have been that network externalities can often lead to multiple equilibria and that converters can have significant impact on equilibrium adoption levels. However, these works only consider static models, that is, they do not incorporate how heterogeneous user decisions lead to the adoption dynamics, and hence, do not model the exact convergence path or identify which one of the several possible equilibria gets realized. Thus, these models are not equipped to study the dynamics of network technology migration.

2.1.2 Model Description

The process of migration from an incumbent to an entrant network technology is governed by the user's adoption decisions. An individual user joins the network technology that offers a higher 'value' in terms of the technology's quality, externality benefits and price. We account for these factors and their effect on technology adoption through a utility function. For each of the two competing technologies, the utility function increases with the technology's intrinsic (stand-alone) quality and the number of other users reachable using it (externality), while it decreases with price. The utility function accounts for user heterogeneity in their evaluation of a technology's intrinsic quality. The network externality benefits that users enjoy from a technology grows in proportional to the number of users that are using the same technology as well as those who are reachable through gateways or converters that their technology deploys. We note that in our model we consider 'technology-level' converters (gateways), *i.e.,* these converters, once deployed, are available to all users of the technology. The model considers that the price (fees) that users pay for subscription to a technology is recurrent because of the service nature of most network technologies.

A user adopts a technology when it provides a utility that is both positive (*i.e.,* satisfies individual rationality constraint) and higher than that of the other technology (*i.e.,* satisfies incentive compatibility constraint). The users continuously re-evaluate their technology choices, and can switch from one technology to another. Since changes in the adoption decision of one user affects the externality benefits of other users as well, they all revisit their adoption decision over time, resulting in the dynamics of technology migration. This process is commonly captured through continuous time models, as in [5]. We study the diffusion process in our work using a similar continuous time model whose solution provides us with the characterization of equilibrium outcomes and system stability.

2.1.3 Summary of Key Results

The analysis of the model reveals a number of interesting behaviors [8, 9, 10]. Some of the main findings are reported here. Firstly, it shows that the adoption process can exhibit multiple steady state equilibrium outcomes; each with a specific range of initial adoption levels of the two technologies. We also find that this behavior may arise both in presence and absence of converters. Secondly, we find that converters can help a technology improve its own standing in the market, and even ensure its dominance while it would have entirely disappeared in the absence of converters. For example, a low-quality but low-cost technology may thwart the success of a better but more expensive competitor by preserving the ability of its users to access adopters of the costlier technology, whose usage would then be limited to a few 'techno-buffs'. Thirdly, we observe a non-intuitive behavior that improving a converter's efficiency can at times be harmful; they can result in lower market share for an individual technology or for both of them. For instance, high market penetration may depend on the combination of a cheap but low-end technology with a high-end but more expensive one to adequately serve the full spectrum of user preferences. A situation where converters allow the better technology to gain market share at the expense of the lesser technology may result in low-end users of that technology dropping out altogether; thereby contributing to a lower overall market penetration. Fourthly, we show that while in the absence of converters, technology migration always converges to stable steady-state equilibrium; this need not be so when technologies deploy converters for compatibility. The presence of converters can create 'boom-and-bust' cycles in which users switch back-and-forth between the two technologies.

The identification of these behaviors allow network providers to realize the potential impact of various economic and technical factors, and thus help devise competition strategies for network migration. Additionally, the knowledge of possible adverse impact of converters that this model reveals is also useful in deciding on policy interventions by regulators.

We have also verified using numerical evaluations that the results obtained from the model are robust to inclusion of additional cost components like switching or learning costs and a broad range of variations in the structure of the user utility functions [10].

2.2 Shared versus Dedicated Network Architectures

2.2.1 Problem Description

The ubiquity and capabilities of the Internet have led to a rapid growth in networked services and applications. This extends well beyond the migration of voice and video onto the Internet, and has the potential to reach areas either traditionally not networked or accessible only through dedicated networks, *e.g.,* health-care, infrastructure monitoring, surveillance, etc. The introduction of such new technologies and services require the network providers to identify the right architecture for their deployment, that is, whether multiple services should share resources on a common network infrastructure, or should each service be offered on a dedicated network of its own. Both these architectural choices have pros and cons. The benefits of a shared infrastructure notwithstanding, combining services with disparate requirements onto a single network also comes at the cost of increased complexity. It often calls for upgrading the network with features required by the new services. This cost scales with overall network size, *i.e.,* is borne by services with no need for the features. It can also introduce complex interactions and the need for tracking and troubleshooting problems of previously little consequences. Therefore, assessing the relative benefits of shared and dedicated networks calls for understanding the trade-off between the economies of scale and scope that sharing allows, and the diseconomies of scope it gives rise to.

A model that helps providers to analyze these trade-offs must capture all the different network deployment and operational cost components, and how these costs are affected by the needs of the services. Additionally, it also has to account for the fact that the actual demand of a new service is initially uncertain, and so the provider has to allocate capacity (resources) in anticipation of the demand. But networks are not the first to face such issues. There is a long tradition of investigating the trade-offs between flexible and dedicated resources and their allocation decisions in the manufacturing systems literature. For example, the Manufacturing Process Flexibility literature has focused on efficient-plant product assignments [14, 15], the effect of process flexibility in handling demand variability [16], and the optimal resource planning and allocation in presence of demand uncertainty [17, 18]. Although the network provider's problem of choosing between shared and separate networks parallels selecting flexible or dedicated manufacturing plants, and making the right capacity allocations, there are key important differences. First, rather than explore the benefits of a flexible (shared) plant (network) in dealing with uncertain demand, our focus is on investigating the impact of various economies and diseconomies of scope in the cost components. A second and more significant difference is that these traditional manufacturing plant models assume that due to large time-lag in building new plants, production cannot be ramped-up rapidly in response to higher than expected demand, whereas in many networks it is quite feasible to increase the capacity on a relatively short time scale, and hence accommodate a portion of the excess demand. This ability to dynamically adjust the network's capacity through resource '*reprovisioning*' is becoming increasingly easy with technological advancements, such as virtualization [20, 19]. The emergence of virtualization technology has made the question of whether to add a new service on an existing network or on a network 'slice' a more practical one. Besides the Internet, this capability is also very common in distributed database networks, cloud computing etc. Even in the traditional manufacturing sector, flexible manufacturing facilities that can be configured nearly "on-demand" are becoming more common. As a result, the earlier manufacturing system models which share some of our structural properties are no longer applicable. They need to be extended to incorporate this new technological ability which not only affects the optimal resource provisioning decisions, but potentially impact the choice of network architecture as well.

2.2.2 Model Description

The most basic setting in which the question of network sharing and all its related economic factors arise is the case of two network services. The model we develop considers the case where a network provider has an existing service that has already been deployed and runs on its existing network. The second service that the provider wants to introduce is a new service with an uncertain demand. We assume that the provider only knows the demand distribution but not the actual demand that will be realized once the service is made available. The provider has to decide whether to deploy this new service alongside the first service on the existing network, or to create another dedicated network for it. Additionally, the provider also needs to decide how much capacity (resources) has to be allocated for the new service given its choice of network architecture. These decisions have to be made prior to the actual realization of the new service's demand. Once the actual demand is known, the provider can increase its capacity if excess demand is realized. Downward adjustments are assumed precluded by contractual obligations. In order to account for different levels of reprovisioning ability, we introduce a *reprovisioning coefficient* that captures the extent to which the provider can recoup excess demand, *i.e.,* a value of 0 implies that all excess demand is lost, while a value of 1 corresponds to an environment were provisioning is not needed as the necessary resources can be secured on-demand and without penalty.

The first step of the model involves identifying the various deployment and operational cost components. We group these costs under the categories of fixed costs, variable costs which grow with the realized demand, and cost of capacity which grows with allocated resources, and use these to formulate the revenue function for each architectural choice. The model we develop is generic enough to allow exploring either economies or diseconomies of scope in the various cost and revenue parameters. We develop a three stage sequential decision process to solve the provider's decision problems. In the first stage, the provider picks one of the architectures. In the second stage, the provider decides on how much capacity to allocate for the new service. The third stage is one where the new service's actual demand finally gets realized and the provider reprovisions resources if necessary to accommodate excess demand. We solve this model by working backwards through the three stages, *i.e.,* given the possible realizations of demand; the corresponding profits are computed in the third stage. Using these demand-profit relationships, the net estimated profit for a given allocated capacity can be computed for each architectural option. This is done in the second stage, where the 'optimal' capacity that generates the highest estimated profit for each of the two architectural options is calculated. Using these optimal capacity decisions, the provider can compare the resulting profits for the two architectures in the first stage, and thus decide to choose the one with a higher profitability. The details of this formulation are available in [21].

2.2.3 Summary of Key Results

In this work, we developed an analytical model that addresses the fundamental issue of network architecture selection. It creates a reasoning framework that includes not only factors like the economic relationships among various cost and revenue components, but also incorporates the impact of demand uncertainties and resource reprovisioning into the decision process. We use this model to show that the extent of (dis)economies in various cost or revenue components can impact the choice of network architecture. The results illustrate the impact that reprovisioning can have on the choice of network solution, and validates the need for models that incorporate such a feature [21]. We also found that the extent of (dis)economies in various network cost components can significantly influence the choice of architecture, and when it does so there is usually a threshold value for these costs where the choice of architecture gets switched. The ability of the model to help the provider quantify these values can guide their selection of more profitable architecture. We are carrying out additional analysis to better illustrate when, why, and to what extent the ability to reprovisioning impacts the choice of shared versus dedicated network architectures.

2.3 Minimalist versus Functionality-rich Network Architecture

2.3.1 Problem Description

Deployment of new network architectures involves deciding on the functionalities that it needs to provide. The degree to which a network architecture is open to outsiders and the range of built-in network functionalities available is critical in enabling the creation and deployment of innovative value-added services. For example, the open Internet architecture allowed many new applications such as email and web to be offered which benefited end-users, service providers and network infrastructure providers. In contrast, there are relatively few service innovations in the closed telephone networks. Although open network architecture has its benefits, there are several issues that the infrastructure providers need to deal with. For example, how to decide what functionalities need to be built into the network and what functionalities should be left to the service providers to develop, should these functionalities be offered *a la carte* or in selected *bundles* for the service providers to choose from, and how should service providers be charged for the functionalities that they use, *i.e.,* should each pay according to the type of functionalities their service use or should all of them be charged a flat fee? These question focus on different aspects of the problem and require slightly different models. The scenario we consider in this research is an open architecture where service providers are allowed to choose the functionalities *a la carte* and pay a flat fee. The decision problem we focus on is for the infrastructure provider to decide whether this open architecture should be functionality-rich or minimalist in its design. It is this decision that determines how the network entices new users and new services, and thus influences how profitable the network is and how much social welfare it generates.

2.3.2 Model Description

The description of the model we present next takes the viewpoint of an infrastructure provider who wants to maximize the profitability of the network it owns. A similar approach may be followed to find the best outcome from the viewpoint of a social planner or policy maker, etc. The model we develop considers that the infrastructure provider creates the network architecture and decides on the amount of functionalities to add. The architecture is an open architecture where any service provider can choose *a la carte* from the available set of built-in functionalities, or may also add their own functionalities, to develop and deploy new services on the network. In return for this, they pay some fixed fee to the infrastructure provider. The deployed services are accessed by end-users (*e.g.,* consumers, enterprises) who also pay a fixed connectivity fee to the infrastructure provider. This scenario thus lends itself to a classical two-sided market model [22, 24, 25] where the infrastructure provider is the owner of the platform which brings together the two sides of the market, end-users and service providers. Such models have been previously used in studying net-neutrality issues [23]. As the regulator of the interaction between the two sides of market, the infrastructure provider also gets to decide the fees that end-users and service providers have to pay so as to maximize the profit earned from the network. But making these decisions requires that the profit or utility functions for the infrastructure provider, service providers, and end-users are all accounted for. Next we discuss how these profit functions are constructed

The profit function for the infrastructure provider is the sum total of revenues earned from service providers and end-users, and from which the cost of incorporating the selected number of functionalities into the network is subtracted. The revenue components are calculated by multiplying the number of services/users with the fees that they pay, whereas the cost component is determined by some exogenous function which monotonically increases with the number of built-in functionalities provided, *i.e.,* creating a functionality-rich architecture is progressively more expensive.

For the service provider's utility function we consider that it increases with the number of end-users connected to the network, while it decreases with the fees charged by the infrastructure provider and the cost of service development. Service development cost is assumed to be an exogenous function that decreases with the number of built-in functionalities in the network. This is because if a lot of functionalities are already incorporated by the infrastructure provider then the service providers don't incur the cost of developing them. While on the other hand, if less built-in functionalities are available then the service providers have to bear these additional costs. Our model allows this cost to be heterogeneous across service providers. As for the service provider's revenues, keeping in line with the latest trend in service offerings, our model assumes that providers generate their revenues from external advertisements rather than end-users. The end-user's utility function grows with the number of service providers on the network and decreases with the fees charged by the infrastructure provider. We also introduce an additional component in the user utility which captures heterogeneity across the users in their valuation of the intrinsic benefits of joining the network. Heterogeneity across both end-users and service providers implies that only those users and service providers who derive a positive utility will join the network. Hence, if the infrastructure provider charges very high fees, then many end-users and service providers may not join the network, which adversely affects the network's profitability. On the hand, if the provider charges too less then the expense incurred in adding the network functionalities may not be recovered. Thus, an optimization problem arises for the infrastructure provider, which we present as a three stage sequential process in our model.

First, the infrastructure provider chooses the number of functionalities to build into the network. This choice determines the cost that the infrastructure provider incurs as well as the cost that service providers will incur in developing and deploying their services. In the second stage, the infrastructure provider chooses the fixed fees to be charged to the services and the end-users. In the third stage, the equilibrium outcome for the two sides of the market is reached and the number of users and service providers that join the network is realized. As commonly done for optimization models with sequential decisions, we solve for these stages backwards. That is, given a set of fees chosen by the infrastructure provider, we solve the third stage equilibrium outcome for the number of service providers and end-users who join the network. Then for a given choice of number of functionalities, we solve the second stage problem by calculating the profit that the infrastructure provider earns from the equilibrium outcome on the two sides of the market, and use this to find the 'optimal' fees which will maximize this profit. Lastly, in the solution to the first stage problem, we find the 'optimal' number of functionalities that gives the highest possible profit for the infrastructure provider. Next we discuss the direction along which we intend to investigate this model.

2.3.3 Future Research Agenda

We propose to explore the model developed to understand how the inclusion of built-in network functionalities affects profitability, and which conditions influence whether an infrastructure provider chooses a functionality-rich architecture or a minimalist design. We also aim to investigate how the optimal number of functionalities included in the network change if a social planner were to decide it instead of an infrastructure provider.

Another direction which we intend to pursue is the impact on service innovation of having an existing version of a functionality that service providers need. Availability of already built-in functionailities may discourage service providers from experimenting with better approaches. Thus, a functionality-rich architecture may even become detrimental to service innovation. We aim to explore this problem in systematic way as a part of the proposed framework.

3. REFERENCES

[1] F. M. Bass. A new product growth for model consumer durables. *Management Science*, 15(5):215–227, 1969.

[2] L.M.B. Cabral. On the adoption of innovation with network externalities. *Mathematical Social Sciences*, 19:299–308, 1990.

[3] 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.

[4] J. Farrell and G. Saloner. Converters, compatibility, and the control of interfaces. *The Journal of Industrial Economics*, 40(1):9–35, 1992.

[5] D. Horsky. A diffusion model incorporating product benefits, price, income and information. *Marketing Science*, 9(4):342–365, 1990.

[6] D. Joseph, N. Shetty, J. Chuang, and I. Stoica. Modeling the adoption of new network architectures. In *CoNEXT'07 Conference*. ACM, 2007.

[7] Katz, Michael L. and Carl Shapiro Network Externalities, Competition, and Compatibility. *American Economic Review*, 75, pp. 424-40, 1985.

[8] Y. Jin, S. Sen, R. Guérin, K. Hosanagar and Zhi-Li Zhang. Dynamics of competition between incumbent and emerging network technologies. In Proc. of *NetEcon'08*, pp.49-54, Seattle, 2008.

[9] S. Sen, Y. Jin, R. Guérin and K. Hosanagar. Modeling the Dynamics of Network Technology Adoption and the Role of Converters. submitted to *Transactions on Networking*. 2009.

[10] S. Sen, Y. Jin, R. Guérin and K. Hosanagar. Modeling the Dynamics of Network Technology Adoption and the Role of Converters. University of Pennsylvania, *Technical Report*. June, 2009. Available at http://repository.upenn.edu/ese_papers/496/.

[11] J. A. Norton and F. M. Bass. A Diffusion Theory Model of Adoption and Substitution for Successive Generations of High-Technology Products. *Management Science*, Vol.33, no.9, pp.1069-1086, 1987.

[12] J. Farrell and G. Saloner. Installed Base and Compatibility: Innovation, Product Preannouncements, and Predation. *American Economic Review*, Vol.76, pp.940-955, 1986.

[13] L.A. Fourt and J.W. Woodlock. Early prediction of market success for grocery products. *Journal of Marketing*, Vol.25, pp.31-38, 1960.

[14] W. C. Jordan and S. C. Graves. Principles of the benefits of Manufacturing Process Flexibility. *Management Science*, Vol.41, no.4, pp.577-594, 1995.

[15] S. C. Graves and B. T. Tomlin. Process flexibility in Supply Chains. *Management Science*, Vol.49, no.7, pp.907-919, 2003.

[16] E. K. Bish, A. Muriel and S. Biller. Managing Flexible Capacity in a Make-to-Order Environment. *Management Science*, Vol.51, no.2, pp.167-180, 2005.

[17] C. H. Fine and R. M. Freund. Optimal Investment in Product-flexible Manufacturing Capacity. *Management Science*, Vol.36, no.4, pp.449-465, 1990.

[18] J. A. Van Mieghem. Investment Strategies for Flexible Resources. *Management Science*, Vol.44, no.8, pp.1071-1078, 1998.

[19] L. Peterson and S. Shenker and J. Turner. Overcoming the Internet Impasse through Virtualization. In Proc. *ACM HotNets-III*, 2004.

[20] J. Touch, Y. Wang, L. Eggert and G. Finn. A Virtual Internet Architecture. In Proc. *ACM FDNA'03*, Karlsruhe, Germany, 2003.

[21] S. Sen, R. Guérin and K. Hosanagar. Shared Versus Separate Networks - The Impact of Reprovisioning. In Proc. ACM *ReArch'09*, Rome, December 2009.

[22] G. Parker and Marshall Van Alstyne. Two-Sided Network Effects: A Theory of Information Product Design. *Management Science*, Vol. 51, No. 10, 2005.

[23] N.Economides and J. Tag. Net Neutrality on the Internet: A Two-sided Market Analysis. *NET Institute*, Working Paper 07-45, 2007.

[24] Jean-Charles Rochet and Jean Tirole. Two-Sided Markets: An Overview, 2004. Available at http://faculty.haas.berkeley.edu/HERMALIN/rochet_tirole.pdf

[25] Jean-Charles Rochet and Jean Tirole. Two-Sided Markets: A Progress Report, 2005. Available at http://idei.fr/doc/wp/2005/2sided_markets.pdf