The Real Options Approach to Standards for Building Network-based Services

Mark Gaynor
Boston University

Scott Bradner, Marco Iansiti and H.T. Kung
Harvard University

In this paper, we show the economic value of flexibility in standards by linking the business concept of market uncertainty to the technical aspects of designing standards and implementing network-based services based on them. We quantify our theory with a options-like model, showing how to maximize overall gain from the market's point of view when creating standards used to build network-based services in highly uncertain markets. We show how network architectures based on standards that promote parallel experimentation such as the end-2-end argument have a greater chance of providing network-based services that meet an uncertain market because of the added value from the ability to innovate. This model provides a framework to understand the tradeoffs between being able to experiment, market uncertainty, business and technical advantages to services with centralized management and how service providers learn from past generations of the service.

1. Introduction

In this paper, we link the business concept of market uncertainty to the technical aspects of designing standards and implementing network-based services built upon these standards. We create a model depicting the economic value of a network-based service to its provider as the service evolves in uncertain markets. Our model is based on real option theory by Baldwin and Clark[1] and builds on our previous paper [2] that placed the evolution of standards in the context of Complex Adaptive Systems (CAS) and shows with an options like model how modularity in protocols increase value. This paper expands our previous work, showing how to get the most from this modular structure.

In order to value flexibility of standards, we study the value of network-based services built from them. We model the value of standards in the context of the value of the services built upon them to those providing the service. Standards that promote services meeting market needs intrinsically have more value than standards not meeting the market.

Our model provides a framework to understand the tradeoffs between market uncertainty and business and/or technical advantages when designing standards used to build services. We show how standards creating environments allowing easy parallel experimentation with services, and then market selection among the experimental services, will likely produce services that match uncertain markets. Furthermore, we show that when market uncertainty is low, standards promoting network-based services with centrally managed structure that are harder to change, but have other business and technical advantages (such as the cost of providing the service) will provide the most overall value. Our model illustrates the value of creating standards that allow a two-tiered structure. The outer layer allowing easy experimentation acts as an incubator, the inner layer allows market selected successful experiments to migrate into the most effective way to manage the chosen service. We show how the Internet and the Phone Network (PSTN) both allow experimentation and migration of successful services to efficient management structure. The main contribution of this paper is the linkage between market uncertainty, and the structure of standards used to build network-based services.

The financial world has proven how valuable options are as a mechanism to limit risk under uncertainty without capping the upside gain. The theory of financial options [3] has been expanded [4] into the theory of real options. This shows that non-linear gains are possible with investments in real world assets such as oil field expansion, investment in information infrastructure, modularity in computer hardware [1], and modularity in IT standards [2].

The end-to-end argument that Clark, Saltzer, and Reed [5] so eloquently explained was one of the most powerful features of the Internet, and a fundamental reason it has flourished. The idea is to keep the network simple, and build any needed complexity into the ends, or edges of the network. The network infrastructure is unaware of e-2-e applications. To some of its believers, the end-2-end argument is a religion, the followers seem like zealots prophesizing the power of innovation with networks that adhere to the end-to-end principle. Our model helps quantify the value of standards allowing end-to-end services. In particular, it illustrates how increasing MU creates more value to end-to-end architecture.

We look back in time at the evolution of the Internet and PSTN in the context of innovation of services to derive our theory. First, is the importance of end-2-end infrastructure within the Internet in fostering innovation to create standards improving the chance of successful network-based services. Next, is how some of the successful services such as email and caching have (or developed) standards that allow these services to migrate into the network. We found a similar structure has developed in the PSTN with popular features in PBXs migrating into core network-based services.. It is important to understand how and why the standards creating the Internet has been so successful in creating services that meet market need to enable current architects of the Internet to continue its successful evolution. Understanding the past will help us understand the implication of current important standards for architectural effectors that break the e-2-e paradigm such as Network Address Translation (NATs), Firewalls, caching, and other middle-ware in regards to the cost of such devices in terms of lost innovation in network-based services.

The next section of this paper promulgates our theory about standards used to build network-based service architecture. In section 3, we explain the value of the best of many experiments. Next, section 4 uses our theory as a baseline to develop mathematics validating our ideas with a real options like model.

2. Theory

The standards and their implementations at different layers of the protocol stack used to build services are the service architecture. Services that appear similar to users may have different architectures. For example, currently in the Internet Engineering Task Force (IETF) there are two very different proposals for standards implementing voice over IP (VoIP). One proposal from the megaco working group in the IETF and co-developed by Study Group 16 in the ITU-T

(H.248), is championed by many phone companies and switch vendors and is built upon a centralized structure of large Internet phone switches, much like the current phone system. The other possibility, Session Initiation Protocol (SIP), allows distributed management in is simplest form, requiring no network intervention, only the two SIP end devices. However, SIP may also be used in a mode with more centralized management via SIP servers. Our theory helps to predict the expected value of services built with each of these types of architectures.

Providing features using SIP compared to H.248 will be very different from the point of view of the service provider. With megaco, the media controller provides the services to the end device with the megaco protocol. A service such as caller-ID is provided by the controller, users can't experiment with different ways to implement caller-ID. In this particular case, the service provider is responsible for the media controller and hence controls the service. However, in SIP, end users have the option of true end-2-end services and can implement caller-ID any way they please, without any cooperation from within the network The users are free to decide what information should be sent, when, with what security, they can experiment, unlike with the megaco model.

We intend our theory to provide a framework to analyze the value of standards in the context of the management structure they allow services to have within a network. It is not intended to give absolute numbers or rules, but to show general relationships between market uncertainty, parallel experimentation with market selection and the benefit of centrally managed service in regards to the value of a set of standards. Our theory is contingency based similar to an argument by Lawrence and Lorsch[6] showing that the best way to manage a business depends on the business. We recognize that the best standards along with how they are implemented changes as conditions vary.

Below we state a formal set of assumptions and a theory based on these assumptions that quantifies the value of standards describing service architectures, such as e-2-e that allow easy simultaneous experimentation. This e-2-e architecture is compared to architectures that provide a more efficient way to manage the service, but where experimentation is harder to accomplish because changes to add new, or enhance existing services must occur within the network core. We believe that if conditions match those set out by our assumptions, then our theory is a reasonable representation of the tradeoffs involved when deciding how to design standards used to deploy network-based services.

Stage one shows that the value a provider of a service receives is random if there is market uncertainty. Next, the theory shows how picking the best from many ways to provide a similar service provides a good chance of achieving a service with features that are a superior market match. In stage two, we expand our theory to account for technical and management advantages from standards allowing centralized type architecture for network-based services. We explain how different architecture can affect the level of experimentation possible when trying to figure out what features a service should have. We present a theory hypothesizing that when the advantage of more centralized service architecture outweighs the benefit of many experiments, standards promoting a centralized management structure may be justified as an initial model. Both stages one and two look at a single generation of a service, in stage three the theory accounts for how services evolve from generation to generation. We hypothesize that at each generation of a service, service providers learn from the previous generation about what will work better for the next generation.

- **Assumption 1:** The market demand for network-based services has a degree of uncertainty. This means that service providers may not accurately predict the value they will receive for providing a service. We denote this market uncertainty as MU.
- **Assumption 2:** Experimentation with services is possible. This experimentation is used to determine what services best matches the current market conditions in the context of what features will be the most popular.
- **Definition 1:** Let X be a random variable denoting the value to the service provider of a single instance of a service.
- **Definition 2:** A service group is a set of service instances, each service instance is available as a choice to the user as a service. Each choice is a different ways to provide the same conceptual service, such as email, to a user.
- **Definition 3:** Let X(i), i = 1 ,,,, n, be a random variable denoting the value to the service provider of providing the i^{th} particular instance of a service within a service group of size n.
- **Theory 1:** E[Max(X(1) ,…, X(n))] >= E(X), that is, the expected value of the maximum of n simultaneous attempts at providing a service instance to some service provider may be far above the expected value. As n increases, the possibility of a truly outstanding market match grows.

One-way to view Theory 1 is in the context of options, having a choice is analogous to having an option. This theory shows the value of standards that promote many experiments such as SIP over standards making experimentation more difficult such H.248.

Next, we state a set of stronger assumptions allowing a deeper theory that considers the management structure of services based on the degree of market uncertainty compared to any advantages to the service provider of providing more centrally managed services. First, we define more precisely how services based on standards have different management structures.

We define a management structure as centralized or distributed based on which position in the network bears the greatest part of management responsible (or expense) for the service. At one extreme are e-2-e type services with distributed management. With true e-2-e services, the end users have total responsibility for the cost because the network provider is not providing any application-specific services. In fact, the network does not even know the service exists, its sole responsibility is the basic transport of data. This scales poorly at the users end since for each new user there is a significant cost incurred by the user to start the service. However, it scales very well for the service provider since they do not need to change anything to support a new service. In contrast, more centralized managed network-based services have better scaling properties from the point of view of the user and worse from the point of view of the service provider.

To better explain how this service management structure taxonomy works we show one examples of a service (email). We present different management structures used to provide email with the Internet set of standards.

Email - Think about the two extreme ways to provide this network-based service, each with a different management structure, but based on the same set of standards. When email started in the Internet, every host wanting email needed to be a mail server as well as an email sender (running something like sendmail). This is a pure end-2-end solution, nothing in the network knows anything about the email, and all management of the service is at the end hosts. This is very flexible and allows innovation, two end users can make a change, test it out, and start using it. However, it is the most expensive way to manage a mail service in the aggregate because each user pays a significant cost. At the other extreme is a single mail server for the world, such as Hotmail could become. This has many business and technical advantages including no transport of data and, the server must only interoperate with itself. However, with a single mail server and service provider, innovation will most likely decrease for several reasons including less service providers to experiment with new features, less competition, and more costly experimentation. This experimentation is more expensive for many reasons including possible feature conflict when adding new features and the cost to protect other users of the centralized system. This example shows how email can be provided with two very different architectures using different management structure: end-2-end compared to a large core based service. The e-2-e architecture allows more innovation while the centralized structure offers lower cost.

The above example shows classification of a service (i.e. email) according to our service management structure taxonomy. It shows how different architectures for a service may either promote experimentation and innovation in providing new features, or efficient management of the service, but with less ability to innovate. It is the standards for these examples that allow what type of management structure is available. With Internet email, both the distributed e-2-e structure (i.e. every desktop running a mail server), or more centralized systems such as Hotmail are possible because of the flexibility of the standards. As it turns out the real world has settled on a mixture including enterprise (or sub-enterprise) level servers using POP or IMAP, and centralized services such as Hotmail.

The theory below helps clarify what it means to allow easy experimentation in a network.

- **Assumption 3:** The function representing the expected value to the service provider of providing the particular instance of a service that best matches a market is non-linear. More experimentation and greater uncertainty increases the expected value. The service provider receives this value by providing the service that best matches the market.
- **Assumption 4:** The less disruptive and less expensive it is to develop and deploy a service, the more experiments there will be. Standards allowing services with e-2-e architecture requiring no network infrastructure change are generally less expensive and less disruptive than more constraining centralized architectures where experimentation requires infrastructure change.

Assumption 4 is very important to standards development. Crafting standards such as H.248 that require changes within the network infrastructure reduce innovation because experimentation is harder. In contrast, standards such as SIP allowing true end-2-end services enable experimentation without permission promote innovation. One example of the value of standards allowing e-2-e services is creating a new HTTP header, one person can implement a new header without asking permission, and start sharing it with friends. If it seems like a good idea, then propose it to the IETF as a proposed standard, giving the market the chance to accept or reject the change. If Tim Berners Lee, the web's creator required permission from a network authority in order to be allowed to experiment with the web, it is unlikely he would have been able to innovate standards for HTTP and HTML, the building blocks of the web.

We next discuss under what conditions the above advantage of experimentation and choice is not enough to outweigh the inefficiencies of managing a distributed service.

In addition to business advantages from economies of scale, centralized architecture has technical advantages for services such as email and voice. One example of this is the email feature being able to access your mail from any computer. With a centralized management structure such as Hotmail this is a very easy feature to implement, but is very hard to implement with more distributed implementations of email.

- **Assumption 5:** For some services there exist business and technical advantages (BTA) pushing service providers to provide services that are more centrally managed. Let this advantage be represented by BTA as defined above.
- **Theory 2:** If (E[Max(X(1) ,…, X(n))] **-** E(X)) < BTA, a service provider should consider providing a service with a more centrally managed architecture. That is, if the advantage of market uncertainty combined with n simultaneous experiments is less than the business and technical advantages of a more centrally managed service, then provide the service with a more centralized management structure.

Theory 2 shows the value of standards promoting experimentation by users Vs experimentation by centralized authorities. One way to view the left side of the equation in Theory 2 is the value of many users experimenting, the right side, is the value of one central manager undertaking a single experiment.

We assume that services evolve over time, in each generation we have n different attempts (experiments) to provide a service offering with a good market match. Thus, each service generation is composed of many service instances from simultaneous experimentation (i.e. a service group), which are the efforts of many different contributors. Our theory incorporates service providers learning from the previous generation of experiments, thus reducing the market uncertainty from generation to generation.

- **Assumption 6:** Services exist for which the market uncertainly decreases in a predictable manner as a function of the service generation.
- **Theory 3:** Service providers are likely to succeed at providing a service with centralized management when the advantage of MU and parallel experimentation does not outweigh BTA.
- **Assumption 7:** Technology changes and alters the space of possible services. One example of this is when PBX's became computerized (SPC) **-** the whole scope of possible features changed.

This means the best standards are flexible to the type of management structure they allow services implemented under them to have. E-2-e type of management structure may work best at first, but then as MU decreases more centralized management is better. Standards should be able to support both types of management structure. Email is one example of this as we discussed before. Standards that enable both distributed and centralized management structure of a similar service provide the flexibility that translates to value over the life of the standards.

One example of this are two competing standards (SIP and H.248) for providing voice on the Internet. SIP allows e-2-e services, while the megaco protocol does not. SIP is also capable of providing the type of centralized management that is mandatory with the megaco protocol. The SIP standard is flexible enough to allow the full spectrum of management structure from centralized to distributed, something the megaco protocol does not allow. Our theory shows the long range value of SIP, at first when MU is high e-2-e use will promote innovation. Later, as MU decreases, SIP proxies and servers will provide the advantages of centralized management provided by the Media Controller in the megaco protocol standard. Initial High MU implies a bad thing for the megaco protocol, the centralized version of Voice over IP because of its lack of flexibility.

- **Theory 4:** If technology changes, market uncertainty may increase.

This theory is foundational in understanding how to design standards used to build services based on not only business and technical advantages, but also on the market uncertainty. It provides a framework to analyze the value of crafting standards in the context of what type of management structure they allow. It shows the tradeoffs between centralized and distributed management structure with respect to market uncertainty, and the number of experimental attempts to provide the service compared to the potential advantage of centrally managing the service. It shows that when a centrally managed service has an advantage from a business and/or technical perspective, the market for the service may still be better met with services that have less central management, but allow more innovative features. It shows the value of flexibility in standards in the context of what management structure the services built upon these standards allow.

The next section provides a framework for a quantitative model, expressing theorem Theory 1 in terms of a classic model in statistics known as maximum or extreme order statistics.

3. Extreme Order Statistics

This section quantifies how choice adds value, and how this value increases as market uncertainty grows. As shown above, it is difficult to match services to markets when market uncertainty is high. To a single service provider, providing a new and innovative service feature is a gamble (similar to rolling the dice): sometimes you win, sometimes you lose, and sometimes it takes a while to determine if you won or lost.

If we assume the value of the new service to its provider has a normally distribution, Figure 1 shows what we expect to happen if we run many parallel service experiments (as discussed in Theory 1 in Section 2). The expected value of a service instance X, is V. Looking at the percentages in Figure 1, we expect that 34% of our experiments will fall between the mean, and +1 standard deviation from it. We expect to need over 769 experiments to find a single service instance that has a value greater than +3 standard deviations from the mean.

U is the maximum of n experiments in Figure 1. This maximum is composed of two different components: the effect of the mean and the offset from the mean. This offset from the mean is itself composed of two parts; the effect of the standard deviation, and the effect of the parallel experimentation. Thus, we can express U in terms of these parts, U = Mean + Q(n)*s.d. Q(n) is defined as how many standard deviations from the mean U is. Intuitively it makes sense that U >= mean. It also follows that the probability of U greatly exceeding the mean increases as n grows.

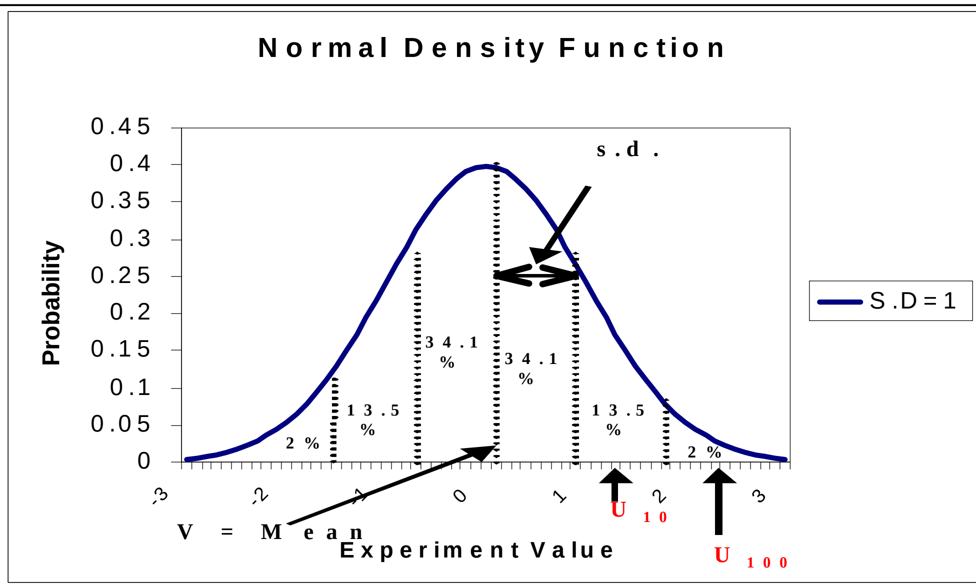

$U = V + ((s.d.) * Q(y))$

Q(y) is the benefit of many parallel experiments

Figure 1. *Best of many experiments*

This model, based on the best of many experiments, can be viewed as real option like since the service provider has the option to pick the experiment that provides a service that best matches the market. This is unlike the case of a single service provide providing a single service choice. This enhanced benefit from having many options is magnified as market uncertainty and the number of service generations increase.

In the next section we use the results above to model the value a service provider is expected to receive by providing a service.

4. Model

In this section, we quantify the above theory by presenting one possible mathematical model. This model is similar in nature to the option based approach by Baldwin and Clark [1] which explains the value of modularity in computer systems design. This model expands on our previous paper [2] about the advantages of modular design in protocol stacks showing how standards allowing flexibility in service architecture may increase the value of modularity.

This model focuses on two main forces affecting the value providers receive for services rendered: the benefit of many parallel experiments combined with market uncertainty pushing services to more distributed management structure, and the efficiencies of centralized management pulling services to centralized architectures. The model is based on the premise that environments providing easy experimentation may not provide the optimal management structure, and environments that are optimized for efficient service management may not be conducive to numerous experiments.

4.1 Modeling a single generation of a service:

Let MU represent the market uncertainty as given by assumption 1. The metric for this may be anything consistent with the metric for BTA defined in assumption 5. X is as defined in Definition 1, the random value to the service provider of providing a particular instance of a service. As before E(X) = V. By the definition of standard deviation (s.d.), s.d.(X) = MU, that is, the standard deviation of the random variable denoting the value of a service to its provider is equal to the market uncertainty. This is because MU is the inability to predict the value of a service, which is just a measure of the variability of the distribution of X.

In our model, we represent the business and technical advantage (BTA) of a centrally managed service relative to its more distributed managed cousin as a cost difference. BTA is the total cost advantage to offering the centrally managed service and may include both management and technical components. BTA is very general, as it must capture all the advantages of centralized management.

Let CP(S) be the cost to provide services with management structure S. E is for end-2-end type services with distributed management, and C is for systems with centralized management structure. This cost is comprehensive and includes both the internal and external components, including internal infrastructure, equipment (including software), and management as discussed in Section 2. It is the total cost to both users and vendors.

In this terminology assumption 5 can be restated as: for some services **CP(E) > CP(C).** It is more expensive to provide services at the end than internal to the network. BTA is defined as:

- **BTA = CP(E) - CP(C)**

VP(S) is the expected value to a service provider with a particular service architecture providing management structure S. This value is just the value received by the provider for providing the service minus the total cost of providing the service.

For internal network-based services we assume that one service is offered, and the value of the service is not restricted to being greater than the mean. From above, the value of a centrally managed service is the expected value of the distribution of received value to the provider, minus the cost of providing the service.

- **Equation 1: VP(C) = V - CP(C)**

For end-based services we assume n service instances in a service group and allow market selection to pick the best outcome as defined in Section 2, theorem 1. As before, Q(n) denotes the value of parallel experimentation, thus the value of the best service with distributed management with the benefit of experimentation in uncertain markets is:

- **Equation 2: VP(E) = V - CP(E) + MU*Q(n)**

The edge is better if VP(E) **-** VP(C) > 0 => MU*Q(n) > CP(E) **-** CP(C), which is equivalent to MU*Q(n) > BTA. So, a service should be provided at the end if:

- **Equation 3: MU*Q(n) > BTA**

This shows Theory 1 from Section 2, as market uncertainty increases end-based services become more attractive because of the enhanced value of experimentation. When the cost differential between providing services with distributed compared to centralized management is less than the benefit gained from high market uncertainty and parallel experimentation. Then the best of the end-based services has greater expected value than a single attempt to provide the service within the network.

Figure 2 shows the relationship between market uncertainty, the business and technical advantage that has been transformed into a cost differential, and the number of experiments run in parallel. Points on the surface represent where

MU equals BTA/Q(n), the points above the surface are the space of services that should be provided end-2-end to take advantage of parallel experiments and market uncertainty. Points below the surface have low enough market uncertainty relative to BTA that centralized architectures should meet market needs. The surface slopes sharply down with regard to the number of experiments showing the value of experimentation. As expected, the space of services benefiting from end-2-end type architectures grows with more experimentation at a decreasing rate. Quickly (around 10 experiments), the rate of decrease levels out showing that the biggest gain from parallel experimentation is from relatively few experiments.

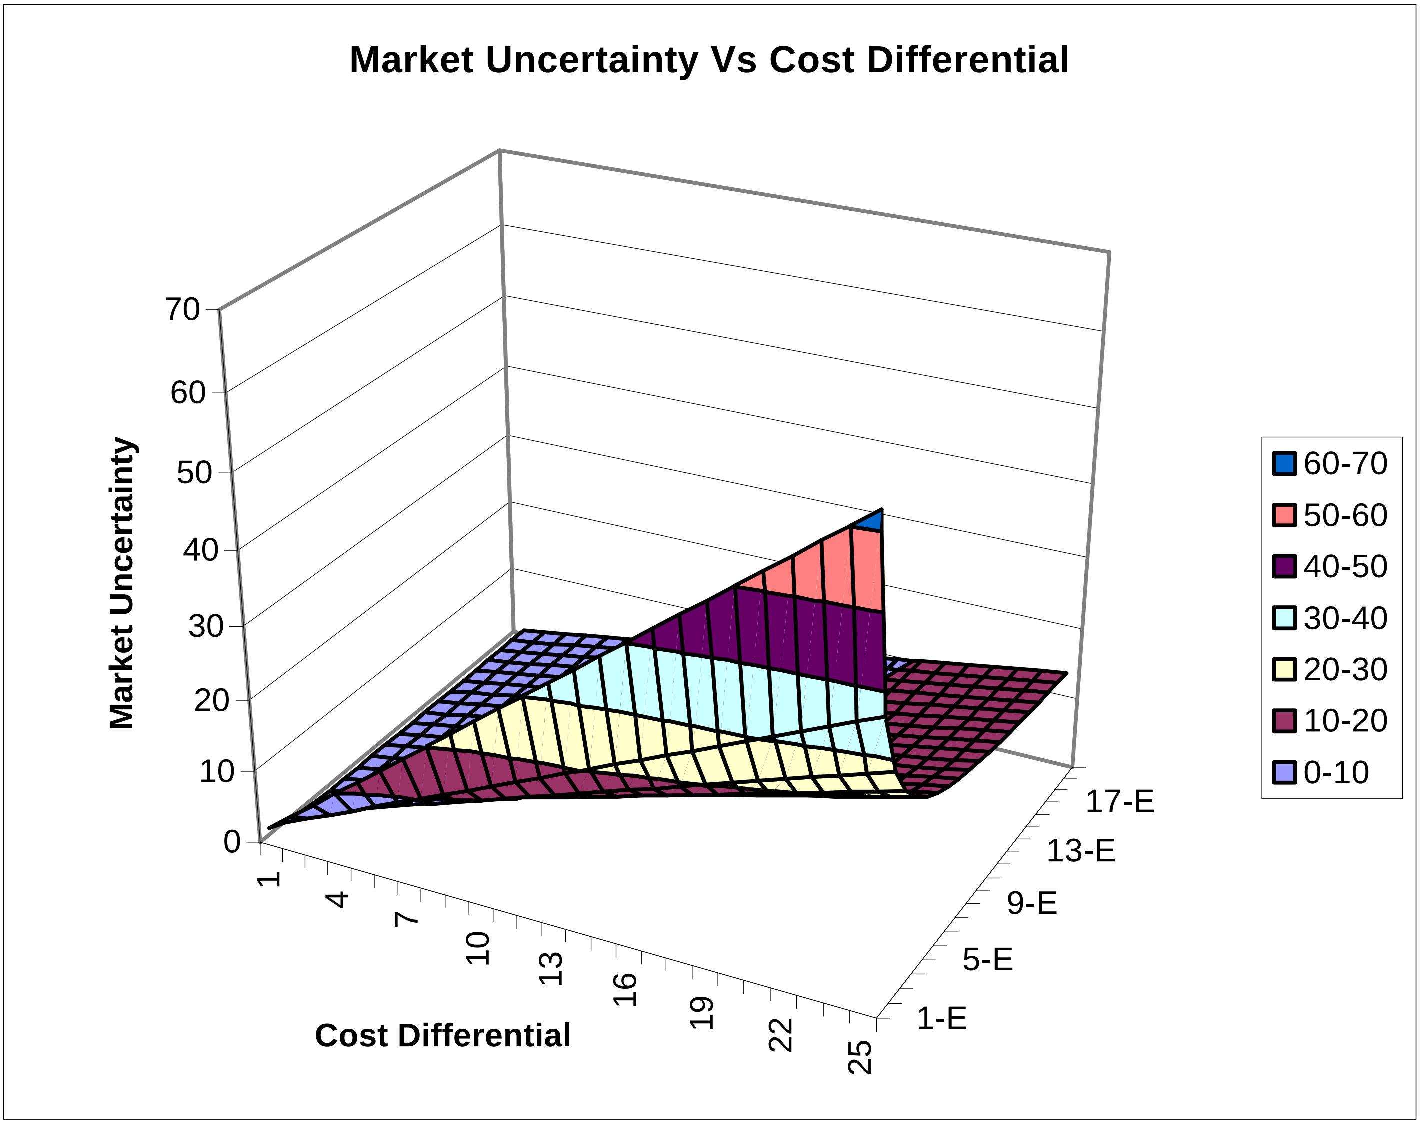

Figure 2. *Simple single generation model*

The above model provides a framework to help understand the relationship between market uncertainty, many parallel experiments, and the advantages of a centrally managed service. It shows how high MU increases the value of standards allowing e-2-e management structure. Next, we expand this basic model showing how services change from generation to generation. In this section, we introduce learning - that is, service providers gain experience from the previous generation about the preferences in the market from watching experiments from other service providers.

The effect of learning is to squash the distribution curve by decreasing the standard deviation (i.e. market uncertainty). Learning has the effect of reducing the benefit of many experiments because each experiment falls within a increasingly narrowing range centered around the mean, thus, many experiments help less and less. To model learning from past generations a function dependent on the generation is introduced to decrease market uncertainty. Let f(generation) be this learning function: f(1) = 1 by definition, there is no decrease in MU at the first generation, and 0 <= f(generation) <= 1, i = 2,3,…,. To simplify the mathematics we assume this learning is symmetric, all service providers learn the same for all experiments run by everybody. Derived from the above single generation model (Equation 1 and 2) the following equations represent this multi-generation model. The value of the first generation is:

- **Equation 3:** $VP_1(C) = V_1 - CP(C)$
- **Equation 4:** $VP_1(E) = V_1 - CP(E) + MU*Q(n)$

.The value of the n^{th} generations is:

- **Equation 5:** $VP_n(C) = VP_{n-1}(C) + V_n$
- **Equation 6**: $VP_n(E) = VP_{n-1}(E) + V_n + f(n) * Mu_n * Q(y_n)$

Solving these difference equations gives:

Equation 7 and 8:

$$VP_n(E) = \sum_{i=1}^{n} V_i - CP(E) \; + MU * Q(y) \sum_{i=1}^{n} f(i)$$

$$VP_n(C) = \sum_{i=1}^{n} V_i - CP(C)$$

This illustrates how the benefit gained from experimentation over all previous generations offsets cost of providing a service with distributed management.

Different types of learning functions invoke dramatically different behavior to the long-term value of evolving end-2-end services. A learning function that diverges (i.e. Harmonic) implies that any advantage of a more centrally managed service can be overcome if the service provider is able to keep evolving the service. However, a convergent learning rate such as any geometric progression strictly limits the advantage gained from market uncertainty and many experiments.

The above equations allow computing value of services at any generation. This allows a similar analysis as shown in Figure 2, but with the number of experiments fixed at 10 (note, the final point on the next two figures represents infinity generations, the limit). Figure 3 shows a divergent harmonic series where the market uncertainty decreases by 1/n at the nth generation. This overcomes any cost advantage (BTA) of centralized management if there is enough time. The surface in Figure 4 shows a very different situation with learning represented as a convergent geometric series. What the series converges to limits the range of cost advantage that experimentation will overcome.

These graphs provide a framework for examining the tradeoffs between market uncertainty, any advantages to a centrally managed service, and how many generations the service is expected to evolve, for a fixed number of experiments. One important question is when, if ever, will the benefits of a centrally managed service overcome the advantage of experimentation to meet the market. Basically, we need to determine at which generation in the evolution of a service the advantages to many experiments is small when compared to the management efficiencies of centralized management structure.

One important question is whether it is better to have fewer generations of a service with more experimentation per generation, or more generations of the service, with less experimentation per generation. With constant MU (i.e. no learning between generations) the decreasing rate of increase of Q(n) implies that more generations with less experimentation is best. However, if MU does decrease, it limits the gain from experimentation making the answer dependent on the rate of decrease.

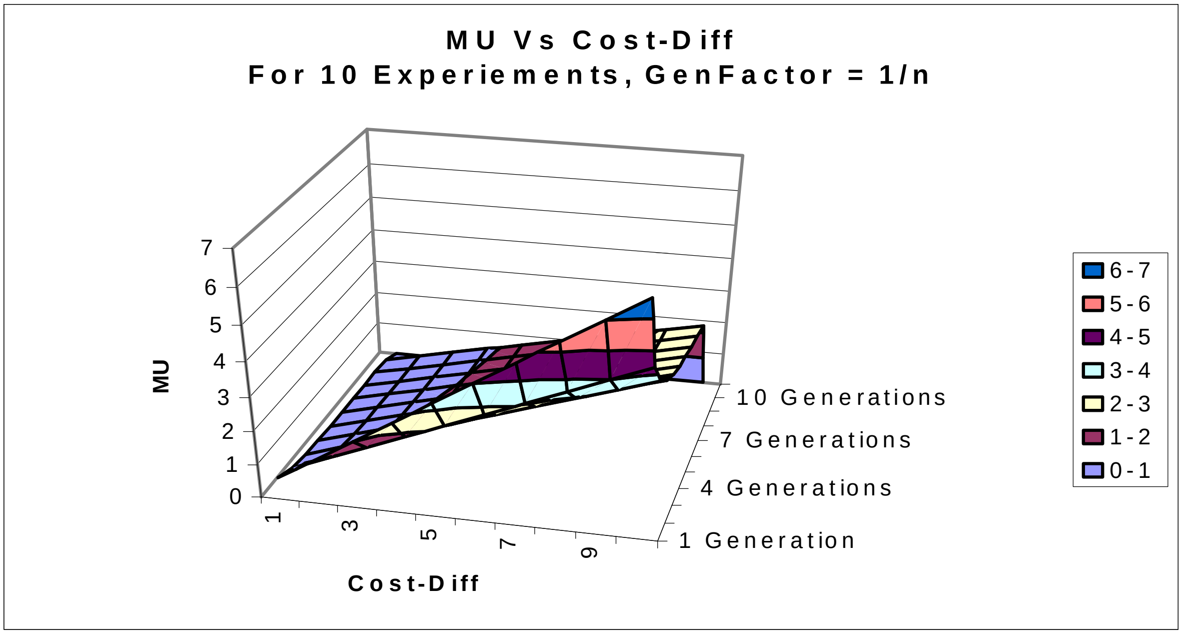

Figure 3. *Harmonic learning (1/n)*

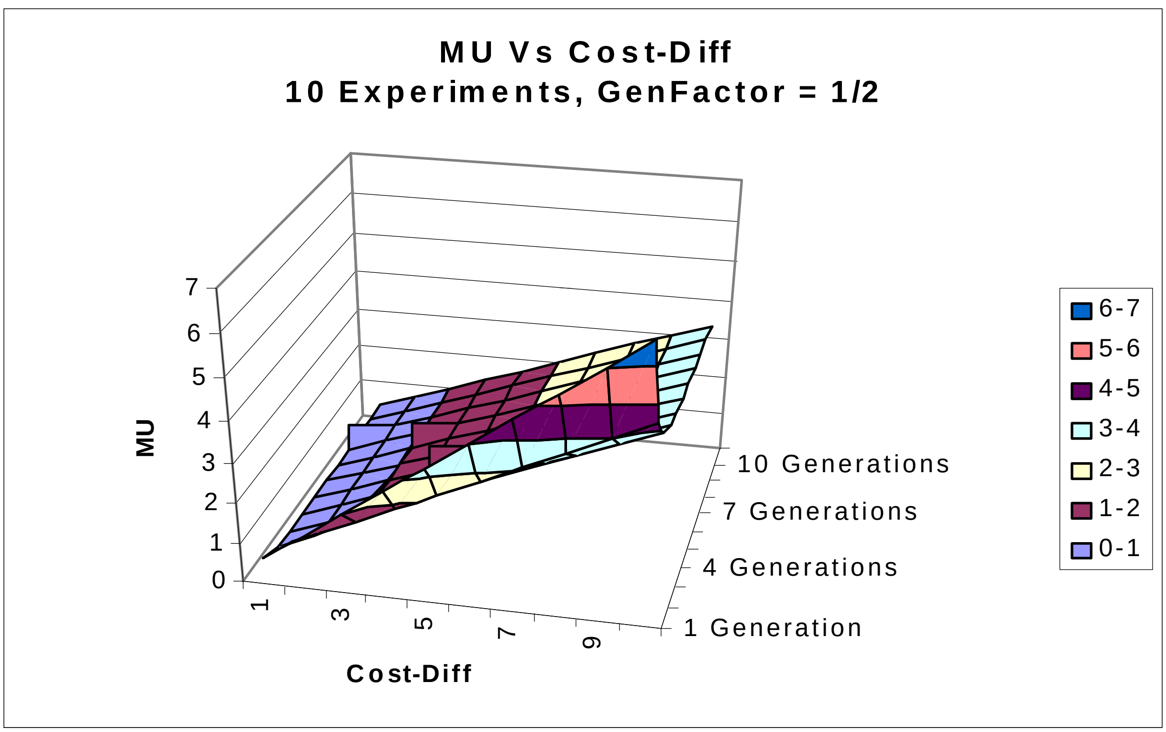

Figure 4. *Geometric Learning (1/2)**n*

The above theory and model provides a framework to better understand the tradeoffs a designer of standards must make when deciding how much flexibility standards should allow in management structure. When the value of a service can not be predicted, our model helps understand the interaction between any possible advantage of a more centrally managed structure, the number of experiments all service providers undertake, and the number of evolutionary generations a service is expected to undergo. It shows how standards promoting flexibility in the choice of management structure they allow have the most value when MU is changing.

While modularity and distributed structure such as e-2-e promote innovation, other factors inhibit the creation of new and better technology. The economic value of network externalities and the related jumping on the bandwagon and user lock-in [7] are examples of innovation inhibitors. These concepts fit nicely into our model. The bandwagoning effect is an indication of lower MU because it shows that customers are deciding. Users flock to technology because others that they respect have adopted it, but it must ultimately meet there needs reasonable well. The economics of network externalities can be used to advantage with new technologies. SIP is a good example of this, by building a SIP/PSTN gateway, SIP users are able to call any phone. This increases the value of SIP because of the vast number of users the SIP caller can reach.

5. Applying the Model

Our research shows the importance of standards allowing a two-tiered structure in Figure 5. On one hand, the ability to provide end-2-end services is necessary to meet user needs in uncertain markets. The ability to try out many different types of services and allowing the market to select the one with the best fit may provide a superior service. However, after understanding consumer needs better, the ability to migrate the services into the network may be necessary to capitalize on the business and technical advantages of centralized management. The outer region gives the power of innovation, while the inner region allows efficient implementations of the best innovations. Standards enabling both e-2-e and centralized management structure promote this type of two-tiered structure.

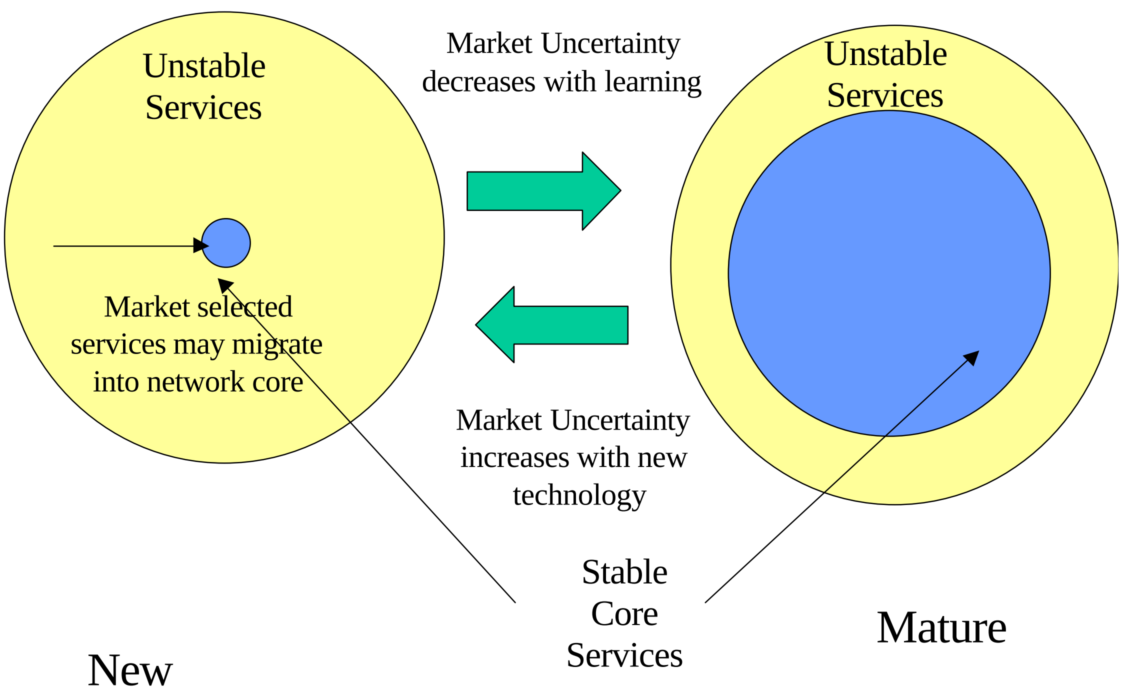

Figure 5. - *Management Diagram*

The Internet supports our two-tiered service industry structure. The Internet started out with very little in the way of internal network-based services, but over time has evolved into a network with more core based services, and middleware. Examples of this are centralized mail services such as Hotmail and Yahoo, middleware solutions to network security such as Radius and Diameter, and QoS. The phone network did the opposite, it started out with no intelligence at the edges, but has evolve into a network that allows more sophisticated Customer Premise Equipment (CPE) (PBX's and smart phones), and even customer configuration of internal network-based services (ISDN). These two examples show that two very different networks have evolved into a similar two-tired structure allowing innovation at the edges, and migration of successful innovations inside the network, thus illustrating the power of this argument.

6. Conclusions

Our theory and model of the economic value of standards allowing flexibility of choice in management structure provides a framework to understand the advantages of experimentation and market uncertainty compared to the financial advantage of services with centralized management architectures. It shows that when users are confused (high market uncertainty) the value of experimentation is high. However, when service providers can predict what services and features will meet market demands, the efficiency of centralized management structure becomes more important than the ability to innovate.

Based on a real options approach our work quantifies the economic value of standards allowing flexibility in choice of management structure. It shows the value of promoting the end-2-end argument when MU is high, and the value of centralized management when MU is low. We have seen that end-2-end services will match markets best and produce the highest value to a service provider when high market uncertainty boosts the benefit of experimentation. However, end-2-end architectures tend to lose some of their attraction as it becomes easy for service providers with more centralized structures to meet market needs at a lower cost.

We found the success of standards promoting a service industry with a two-tiered structure. The high end of the market being better served by providers that have the ability to experiment and innovate the cutting edge features, while the normal users are often well served by service providers with more centralized structures. The success of centralized services depends on the market uncertainty being low enough relative to the economic advantage of centralization that a high percentage of the market can be satisfied.

This work links flexibility in standards to market uncertainty. Understanding how and why successful services have evolved in the PTSN and Internet is important to continue the innovation. This is particularly important in the age of convergence of data and voice services.

7. References

[1] Baldwin, Carliss, and Clark, Kim, *Design Rules: The Power of Modularity*. Cambridge, Ma: MIT Press 1999.

[2] Gaynor, Mark, and Bradner, Scott The Real Options Approach to Standardization, *Proceedings of Hawaii International Conference On Systems Science* 2001.

[3] Brealey, Richard, and Myers, Stewart, *Principles of Corporate Finance*, McGraw-Hill, 1996

[4] Amram, Martha, and Kulatilaka, Nalin, *Real Options, Managing Strategic Investment in an Uncertain World*, Harvard Business

[5] Saltzer, Jerome, and David, Reed, and Clark, David, End-To-End Arguments in system design, *ACM Transactions in Computer Systems* 2, 4, Nov 1984 p 277-288.

[6] Paul Lawrence, and Jay Lorsch, *Organization and Environment*, Harvard Business School Press, 1986

[7] Shapiro, Carl, and Varian, Hal, *Information Rules - A strategic Guide to the Network Economy*. Harvard Business School Press 1999.