The Real Options Approach to Network-based Service Architecture

Mark Gaynor (gaynor@eecs.harvard.edu), Scott Bradner (sob@harvard.edu), Marco Iansiti (miansiti@hbs.edu), and H.T. Kung (htk@eecs.harvard.edu)

ABSTRACT

In this paper, we link the business concept of market uncertainty to the technical aspects of designing and implementing network-based services. We quantify our theory with a options like model, showing how to maximize overall gain from the market's point of view when providing network-based services in highly uncertain markets. We show how network architectures 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 and implementing network-based services with a model showing the 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 how modularity in protocols increase value. This paper expands on this showing how to get the most from this modular structure.

Our model provides a framework to understand the tradeoffs between market uncertainty and business and/or technical advantages to building services with centralized management structure. We show how 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 how and when certain types of network-based services can take advantage of centrally managed services that are harder to change, but have other business and technical advantages (such as cost). Our model shows the value of 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 validate our model by showing how the Internet and the Phone Network (PSTN) both display attributes of this two-tiered structure. The main contribution of this paper is the linkage between market uncertainty, and how 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] showing how non-linear gains are possible with investments in real world assets such as oil field expansion, investment in information infrastructure, and modularity in computer hardware [1].

The end-to-end argument (stupid networks) that Clark, Saltzer, and Reed [5] so eloquently explained was one of the most powerful features of the Internet, and a fundamental reason it flourished. 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 the end-to-end argument in at least one area. In particular, we examine what classes of network-based services, in what conditions, can best take advantage of end-to-end architecture.

We looked 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 successful network-based services. Next, is how some of the successful services such as email and caching have migrated 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 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 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 of 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 an real options like model. The final section discusses how to apply our model to services in the real world.

2. Theory:

The service architecture explains the protocols that can be used at the different network layers to implement the service. Services that appear similar to users may have different architectures. For example, currently in the IETF there are two very different proposals to implement voice over IP (VoIP). One proposal, megaco, championed by many phone companies and switch vendors is built upon a centralized structure of large Internet phone switches, much like the current phone system. The other possibility, SIP, is managed as a more distributed system, in its simplest form, it requires no network intervention, only the SIP end devices. Our theory helps to predict the expected value of services that will be built with each type of architecture.

Providing features using SIP Vs megaco 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.

Our theory is intended to provide a framework to analyze how and when to provide services 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 particular service. 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 architecture for building services changes as conditions vary.

Below we state a formal set of assumptions and a theory based on these assumptions that quantifies the value of 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 must occur within the network core to add new, or enhance existing services. 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 network-based services.

Stage one ([A,T,D]1.*) shows that the value a provider of a service receives is random because of 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 ([A,T]2.*), we expand our theory to account for technical and management advantages from 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, 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 ([A,T]3.*), 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.

- **A1.1:** 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.
- **A1.2:** 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.
- **D1.1:** Let X be a random variable denoting the value to the service provider of a single instance of a service.
- **D1.2:** A service group is a set of service instances, each service instance is available as a choice to the user as a service.
- **D1.3:** 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.
- **T1:** 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 T1 is in the context of options, having a choice is analogous to having an option. In the case of services, this means users have the option to pick from many services within a service group, the service that best matches their needs. Our theory follows a result of options theory: the value of a portfolio of individual options is worth more than an option on the portfolio as a whole.

Next, we state a set of stronger assumptions allowing a deeper theory that considers the location of services within the network based on the degree of market uncertainty Vs any advantages to the service provider of providing more centrally managed services. First, we define more precisely how a service is managed in the context of how the management is split between the user and service provider.

We define the location of a service within the network as the position in the network that bears the greatest part of management responsible (or expense) for the service. At one extreme, with true e-2-e services the end users have total responsibility for the cost because the network provider is not responsibility for end level application services and, in fact, the network does not even know the service exists, the network's sole responsibility is basic transport of data. This scales poorly since for each new user there is a significant cost incurred by the user to start the service. In contrast, more centralized managed network-based services may have better scaling properties. The marginal cost to users of a new service might be close to zero since the service may be compatible with the users current hardware/software configuration (i.e. many web based services).

To better explain how this service taxonomy works we show two examples of services (email and voice services), and present different ways to provide each of these services, but with service architectures that are different in respect to the locality of management.

Email - Think about the two extreme ways to provide this network-based service. When email started in the Internet every host wanting email needed to be a mail server as well as a 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 host which the user is on and the host to which the mail is directed. 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. At the other extreme is a single universal mail server for the world. 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 example shows how email can be provided with two very different architectures using different locations in the network: end-2-end Vs a large core based service. The e-2-e architecture allows more innovation while the centralized structure offers lower cost.

Next, consider how features for dial-tone service, basic voice services, and advanced voice services are provided. Most businesses have two choices for voice services: PBX (end based solution), or Centrex (a centralized PBX running in the Central Office (CO) of the phone company). PBX's in general have more cutting edge features than Centrex service because experimentation is more expensive in CO switches for many reasons including feature conflict and extraordinarily rigorous testing, and the greater number of PBX vendors than CO switch manufactures creates more opportunity for experimentation in PBXs. The PBX solution is locally managed giving flexibility on site while the Centrex service is centrally managed, providing the advantages of a turn-key service. The rational behind Centrex was to save money and provide the same service as a PBX without the inconvenience to the customer of on-site equipment[7], with many popular Centrex features starting out on the PBX. This example shows how even in the historically centrally controlled PSTN different methodologies exist for providing voice services.

The above examples show classification of services (i.e. email and voice) according to our service location taxonomy. It shows how different architectures for services may either promote experimentation and innovation in providing new features, or efficient management of the service, but with less ability to innovate. The theory below helps clarify what it means to allow easy experimentation in a network.

- **A2.1:** 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.
- **A2.2:** The less disruptive and less expensive it is to develop and deploy a service, the greater number of experiments there will be attempting to provide a service that matches the market. Services such as e-2-e requiring no network infrastructure change are generally less expensive and less disruptive.

Assumption A.2.2 is very important to innovation. Changes within the network infrastructure require permission of those in control of the network. With true end-2-end services, new services can be added without permission. For example, one person can implement a new HTTP header without asking and start sharing it with friends. If it seems like a good idea it can be proposed to the IETF, 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 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. Centralized "core" email systems can implement distributed shared access of email easily. Centrex based PBX-like service is technically better suited to providing service in separate locations (as long as it is in the same LATA) than PBXs.

While the above points out that some services have advantage over centralized structures, there are services such as DNS where this is not true. The DNS service is mature, and has not evolved into a centralized managed service. DNS has become more centralized in many environments, but at the institutional level. In the case of DNS, analysis of the data management cost (in particular where the data is sourced and maintained) shows that centralization is only effective to a point.

- **A2.3:** 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.
- **A2.4:** There are services for which market uncertainty is low relative to BTA and this uncertainty will remain low with high confidence.
- **T2:** 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 that service can be provided further within the network.

We assume that services evolve over time, in each generation we have n different attempts (experiments) to provide a service 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.

- **A3.1:** Services exist for which the market uncertainly decreases in a predictable manner as a function of the service generation.
- **T3.1:** Service providers are likely to succeed at providing a service with centralized when the advantage of MU and parallel experimentation do not outweigh BTA.
- **A3.2:** 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.

From the 1920's when PBXs came into use, until the mid 70's PBX were Electro mechanical devices. By the 70's the market and technology was very mature, but in the mid-70's PBXs adopted the Stored Program Control (SPC) architecture that was popular with the new generation of CO office switches. This proved to be a major paradigm shift due to the ease of adding new features in software in comparison to the expense of enhancing the older mechanical based PBXs. This shift drastically increased market uncertainty since the scope of possible features was expanded, and the ease of experimentation with these new features allowed vendors to experiment and differentiate their products, allowing users to select which of the choices met their newly evolved needs best. Another example of this is the shift from dumb to smart phones. With Internet technology such as mini-browsers the definition of voice service is changing. This High MU should imply a bad thing for megaco, the centralized version of Voice over IP.

- **T3.2:** If technology changes, market uncertainty may increase.

This theory is foundational in understanding how to design services based not only on business and technical advantages, but also on how well the market is understood. It provides a framework to analyze the management of service architecture with respect to market uncertainty, and the number of experimental attempts to provide the service Vs 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.

The next section provides a framework for a quantitative model showing how theorem T1.1 is expressed 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.

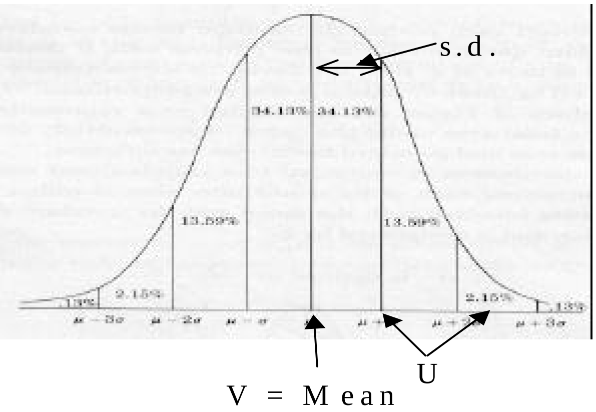

Figure 1 **Best of many experiments**

If we assume the value of the new service to its provider is normally distributed, Figure 1 shows what we expect to happen if we run many parallel service experiments (as discussed in T1.1 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.

The maximum of n experiments is denoted as U 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.

This model, based on the best of many experiments, can be viewed as real option like since the users have the option to pick the experiment that provides a service that best matches their needs, 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 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 A1.1. The metric for this may be anything consistent with the metric for BTA defined in assumption A2.4. X is as defined in D1.1, 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(L) be the cost to provide services with management at location L. E is for end-2-end type services and I is internal. 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 A2.4 can be restated as: for some services **CP(E) > CP(I).** It is more expensive to provide services at the end than internal to the network. BTA is defined as:

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

Let VP(L) be the expected value to a service provider with a particular service architecture providing a service at L. 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 an internal service is the expected value of the distribution of received value to the provider, minus the cost of providing the service.

- **E5.1: VP(I) = V - CP(I)**

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 T1.1. As before, Q(n) denotes the value of parallel experimentation, thus the value of the best service at the edge with the benefit of experimentation in uncertain markets:

- **E5.2: $VP(E) = V - CP(E) + MU*Q(n)$**

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

- **E5.3: $MU*Q(n) > BTA$**

This shows T1.1 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 at the network's end Vs internal is less than the benefit gained from high market uncertainty and parallel experimentation, the best of the end-based service 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.

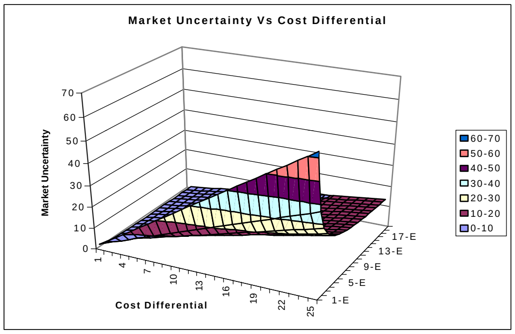

Figure 2 Simple single generation model

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. Next, we expand this basic model to show how services change from generation to generation in their search for a better match in uncertain markets. In this section, we introduce learning - that is, service providers gain experience from the previous generation about the preferences in the market.

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 $0 <= f(i) <= 1$, $i = 2,3,\ldots,$. 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 E5.1 and E5.2) the following equations represent this multi-generation model:

- **E5.4.1:** $VP_1(I) = V_1 - CP(I)$
- **E5.4.2:** $VP_1(E) = V_1 - CP(E) + MU*Q(n)$

.The value of the nth generations is:

- **E5.5.1:** $VP_n(I) = VP_{n-1}(I) + V_n$
- **E5.5.2**: $VP_n(E) = VP_{n-1}(E) + V_n + f(n) * Mu_n * Q(y_n)$

These difference equations can be solved giving equations:

E5.6.1 and 5.6.2:

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

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

This illustrates how the greater cost of providing a service at the edge is offset by the sum of benefits gained over all previous generations.

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 set to 10 (note, the final point one the next two figures represents infinity generations, the limit). A divergent harmonic is shown in Figure 3 the market uncertainty is decreased by 1/n at the nth generation, any cost advantage (BTA) of centralized services can be overcome. The surface in Figure 4 shows a very different situation where learning is represented as a convergent geometric series. The range of cost advantages that can be overcome is limited by what the series converges to.

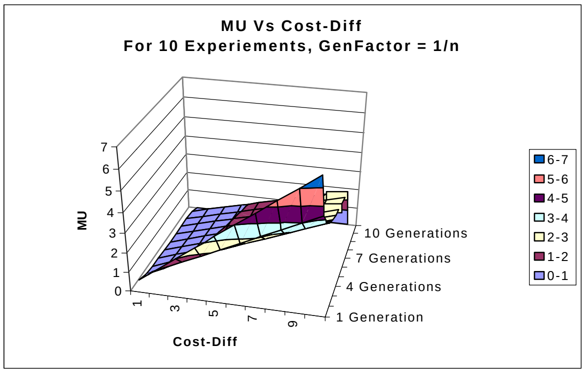

Figure 3 Harmonic learning (1/n)

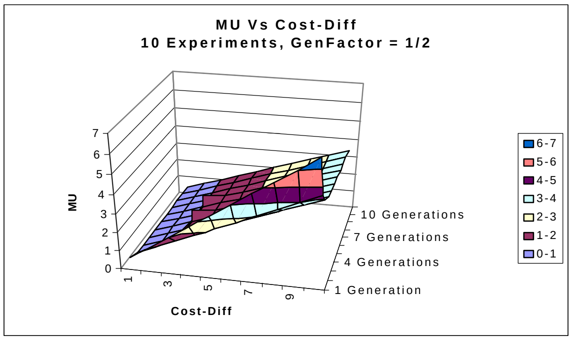

Figure 4 Geometric Learning (1/2)**n

The above 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 network-based services. From the equations above, we can find both lower and upper boundaries as to when a single attempt to provide a centralized network-based services is likely to be successful in meeting market demands.

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.

The above theory and model provides a framework to better understand the tradeoffs a service provider must make when deciding how to design a service architecture to provide a service. When a service provider can not predict a market value for a service, 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.

5. Applying the Model

Our research shows the importance of having 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, once consumer needs are better understood, 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 the best innovations to be implemented with the most efficient architecture.

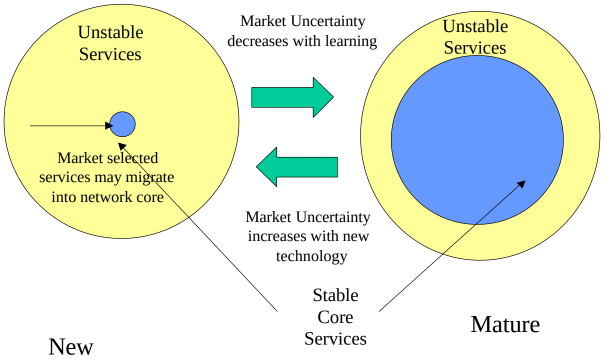

Figure 5 - Management Diagram

The history of the PBX Vs Centrex service illustrates this structure. PBXs are a way to test, and grow services, that if successful can migrate into Centrex services, such as voice mail, and ACDs. Today, many customers that want "cutting" edge services, seek out PBXs, and may even use standalone applications such as Octel's voice mail, or Rockwell's ACD. However, many users (some of them as sophisticated as Harvard or Berkley) have used Centrex since the late 80's and found it meets their needs well. These customers found that the lag between successful features migrating from PBXs, into services within the telco's switch has been less than expected.

The Internet also support 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 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 services in regard to their design and implementation structure provides a framework to understand the advantages of experimentation and market uncertainty Vs 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, but when service providers can predict what services and features will meet market demands, the management structure of the service becomes more important than the ability to innovate.

Our work is an real options based approach to quantify the end-2-end argument, showing the value of end-2-end services because of more innovation. 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 that the service industry is two-tiered in 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. We found that the most successful service in the Internet (email) has this structure, as well as basic and advanced voice services provided by the PSTN.

This work links innovation in network design to architecture of service design. 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, To be published in the *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] Goeller, Lee, Ten Common Misconceptions About Centrex, *Business Communications Review*, May- June 1986