# A Real Options Framework to Value Network, Protocol, and Service Architecture

Mark Gaynor
Boston University
mgaynor@bu.edu

Scott Bradner
Harvard University
sob@harvard.edu

## ABSTRACT
This paper proposes a real options framework for evaluating architectural choices and the economic value of these alternative choices of networks, protocols, and services. Using proven financial techniques of real options, our model explores the value of distributed architecture compared to the benefits of centralized control. Voice and email case studies that agree with our theory and model are presented. We apply our model to illustrate the value of end-to-end structure, why SIP-based VoIP is winning, and the value of open garden service business models allowing third parties to provide network services/applications. This work illustrates the potential of real options to help quantify the economic value of network, protocol, and service architectures.

## Categories and Subject Descriptors
C.2.1 [**Network Architecture and Design**]: Distributed networks, Centralized networks.

## General Terms
Management, Design, Economics, Experimentation, Standardization.

## Keywords
Network architecture, real options, network services, end-to-end.

## 1. INTRODUCTION
The contribution of this paper is a real options based framework linking the value of flexibility (i.e., the ability to innovate) to market uncertainty. Our model illustrates the tradeoff between: (1) who has the ability to experiment, (2) the degree of difficulty of the experimentation, (3) market uncertainty, (4) the business and technical advantages of a more efficient centralized structure, and (5) how architecture evolves over time. This model is useful in evaluating network architecture principles such as the end-to-end argument [1], comparing network infrastructure variations (such as circuit Vs packet switching), protocols within a particular category (such as SIP Vs Megaco/H.248 for voice over IP), and service architectures built with similar protocols (such as centralized Vs distributed music distribution systems). This framework shows that when market uncertainty is high, network, protocol, and service architectures that foster experimentation at the network's edge creates a potential for a greater value for a service provider than does central administration of network activities.

Our framework is useful to designers and managers making decisions about network, protocol, and service architectures. Having imperfect information about what users prefer and how much services are worth to these users, designers must make many decisions about distributed compared to centralized architecture, converged networks, and open Vs walled garden service architecture. One application of this theory illustrates the importance of the end-to-end argument by explaining how high market uncertainty increases the value of open end-to-end structure to service and software providers. Our model illustrates why some current network architecture, such as Network Address Translation (NAT) and firewalls, that break end-to-end ideas at a fundamental level will have devastating consequences to the opportunities for service providers to profit from tomorrow's Internet. The fact of data, voice, and video convergence is clear, but the architecture of protocols in the converged network (for example, Voice Over IP (VoIP) protocols) is still uncertain. Our framework helps evaluate open standards such as the IETF's Megaco (a.k.a. the ITU-T's H.248, which is an example of an architecture that assumes a centrally managed structure), compared to the IETF's SIP, (which is an example of an architecture that allows both centralized and end-to-end management structures), compared to proprietary protocols such as Cisco's Skinny. Broadband Internet connectivity is coming, but the business models that carriers will use is still not clear: How will services be provided over broadband Internet connections? Will centralized service providers control what content and services users are able to access in the name of increasing the service provider's income? Or will a more open garden business model, where users have many choices for content and services, prevail?

Decisions made today will affect the nature of innovation in tomorrow's network infrastructure; the choices that will create the most overall value will encourage network infrastructure that promotes innovation, yet also allow for efficient scalable services. Our model is useful for evaluating architectural choices and the economic value of these alternative choices that designers have.

## 2. DISTRIBUTED VS CENTRALIZED
The management structure of a network, protocol, or service/application not only affects the cost and complexity of experimentation, but also the range of who can participate in carrying out experiments, which is discussed in detail in [2][3][4][5]. A distributed structure promotes innovation by encouraging cost-effective experimentation with new ways to provide services that anybody (including end users) can perform. This is in contrast to systems with centralized structure that wind up discouraging experiments because of

cost and complexity of experimentation, or because of the reluctance of the network or service manager to allow it.

An example of why experimentation is promoted in a distributed architecture and inhibited in a centralized architecture is depicted in two possible different structures for a music service (See Figure 1). In the distributed architecture, Bob and Alice can use completely incompatible systems because each system only has to deal with its own group of users. Furthermore, in a distributed architecture, each user is able to develop his or her own new devices or protocols. This is very different than the centralized model, which utilizes a centralized music server. This architecture may be very efficient to operate, but it is not as flexible since the centralized service provider must coordinate any protocol changes and new devices with all of its users.

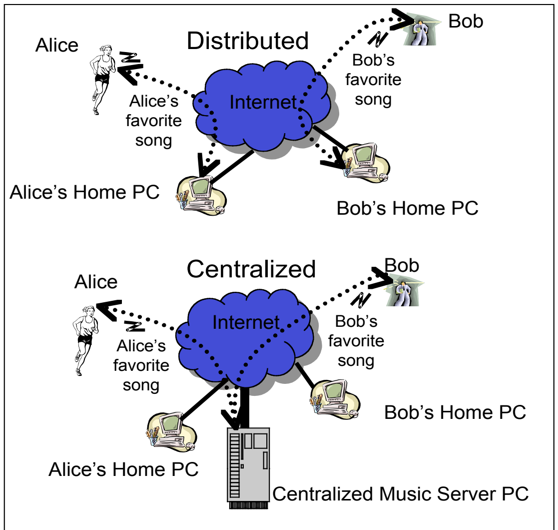


**Figure 1** Distributed Vs Centralized Music Server

# 3. VALUE OF EXPERIMENTATION

When a new type of product or service is first introduced, equipment vendors or service providers may not understand what features potential users will want since neither the users or the providers have any experience with the concept. It is possible that the first company to introduce the product or service will get it right, but it's far more likely that they will not and the field will be open for other vendors to make their own attempts to discover just what the customers want. We will call these attempts to find what the customer wants "experiments."

Vendors experiment with different applications, as well as with similar applications that have different feature sets. Each experiment is seen by the user as a product and is an attempt to meet an uncertain market. The economic value of experimentation links to market uncertainty by definition: uncertainty is the inability of the experimenter to predict the value of the experiment (i.e., predict what the customer wants). When uncertainty is zero, the outcome of each attempt to meet the market is known with perfect accuracy: the market match is perfect every time because user preferences are completely understood, and this makes the products/services a commodity where no one makes much money. In low market uncertainty, it's hard for any vendor to win big because competition becomes price based, but when uncertainty is high, successful products are likely to generate huge financial success because they can capture a large part of the market with a unique product/service.

## 3.1 What is an Experiment

Experiments are vendors, service providers, entrepreneurs, and users trying to meet a perceived customer need with a new architecture, protocol structure, service/application, or a particular feature set for a service/application. Experiments happen at many levels from network architectures to particular feature sets built on accepted standards. At one time, the OSI set of protocols and the Internet Protocols were two experiments at the network architecture layer. Internet Protocol standards became the dominant design, which opened up a new layer of experimentation: protocol structure within the Internet framework, for example Tim Bernners Lee's web experiment. A current example of experimentation is with protocols used to provide VoIP such as SIP, Megaco, Skinny, and H.323. Within each protocol, such as Megaco or SIP, there is experimentation on features within each protocol. There is also experimentation at the service/application layer. If SIP is being used, what should you do with it? Chat to friends over your Internet connection, or maybe integration of your voice mail and email, or … What features should these applications have, and how should they work? The point is nobody really knows, but many service providers/vendors/users are trying many ideas–each one of these attempts is an experiment.

There are many other historic examples of experimentation–some successful, some not. The PBX and its many standard features emerged from experimentation with thousands of features from vendors when the PBX design became computer based [5]. A standard feature set emerged, and successful PBX features such as caller-ID, speed dialing, and voicemail have been retrofitted into Centrex (a carrier-based PBX service). Not all experiments are successful: for example, when the PBX vendors in the 80's experimented with offering a switched data service using their TDM framework, it failed in the marketplace. There are also examples of failed Internet standards: Privacy Enhanced Mail (PEM) and Simple Network Management Protocol v2 were both unsuccessful experiments. PEM failed because required international public key infrastructure never materialized, and customers felt local security infrastructures did not offer enough benefits for customer use. SNMPv2 failed because customers felt that it was too complicated. When market uncertainty is high, we expect a few experiments to be highly successful and expect some others to be spectacular failures.

In the lists of experiments above, there is an important thread: in each case, market selection by users determines the value of these experiments by choosing what matched his or her own needs best, at a price they are willing to pay. It is this experimentation, along with market selection of the most successful experiments that creates the options value of allowing users to have choice.

## 3.2 Best of Many Experiments

If the potential value of each experiment falls within a normal distribution, then Figure 2 shows what we expect to happen when attempting several parallel experiments. It illustrates the probability of experiments being a particular distance from the mean. $V = E(X)$ denotes the expected value (mean) of a particular experiment, U(n) denotes the value of the best of n parallel experiments. Looking at the percentages in Figure 2, we expect that 34 percent of the experiments will fall between the mean and +1 standard deviation from it, and so on. This illustrates that when the standard deviation (s.d) is high and your experiment is the most successful at meeting the market, then it is likely you will win big by capturing most of the market because your product better meets users wants in a market where products are differentiated by features. This chance to win big will induce many vendors to experiment, which increases the probability of a superior match and big win for the lucky (or smart) vendor or service provider.

Figure 2 illustrates $U(10)$ and U(100), the expected maximum from 10/100 experiments. This maximum is composed of two different components: the distribution mean (V) and the offset from the mean. This offset from the mean is composed of two parts: the effect of the standard deviation and the effect of the parallel experimentation. Thus, the best of many experiments ($U(n)$) can be broken into these parts: $U(n) = V + Q(n)*s.d.$ $Q(n)$ [6] measures how many standard deviations from the mean *the best of n experiments* is, and the market uncertainty is the scaling factor. The probability *of* the best experiment greatly exceeding the mean increases as the number of experiments grows or the standard deviation increases.

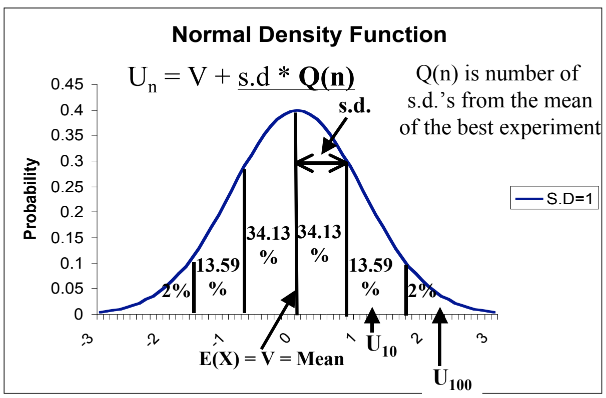


Figure 2 How Uncertainty Changes the Value of Experimentation

The best of these experiments has a value that grows further from the mean, but at a decreasing rate. As market uncertainty increases, so does the gain from experimentation and thus, the potential for profit. It demonstrates that with low market uncertainty, even a large number of experiments have low value because the scaling factor overrides the benefit from experimentation.

In other words, if everyone knows exactly what the customers (market) want, all vendors will make essentially the same product and it will be a mostly undifferentiated commodity, which implies that no vendor will make much money because they will split the business and compete on price. On the other hand, if no one has a clue of what the customer actually wants, the vendor that guesses correctly will capture the market (and customers) all to themselves and will be able to charge whatever the market will bear. The chance to win big induces investment, which increases experimentation–a wining situation for users (and the winning entrepreneur).

# 4. A REAL OPTIONS FRAMEWORK

The above statistical argument links MU to network architecture and gives a good basis for some assumptions, a theory about the architecture of networks, protocols, and the services built with this infrastructure, as well as a supporting model, which we present below.

## 4.1 Model Assumptions

Before presenting an intuitive model based on the mathematics of real options, we lay out some simple assumption and a few basic theories about how market uncertainty links to network infrastructure architecture.

- There is uncertainty about users preferences; we call this market uncertainty. This means that network owners, protocol designers, and service providers cannot accurately predict the value to their users of what they are building/providing. This market uncertainty is denoted as MU.
- Experimentation with network infrastructure, protocol structure, and services is possible. Different types of architecture allow different types of experimentation. Market selection picks the best from the many trials. This is how the experimentation is evaluated to determine what matches the market best as well as its value.

The value of the maximum of $n$ simultaneous experiments is greater than the expected value for any particular experiment. From above, we know that as the number of experiments or the standard deviation (i.e. market uncertainty) increases, so does the difference between this maximum of $n$ experiments compared to average value or a single experiment. With high market uncertainty and many experiments, the possibility of a truly outstanding market match grows.

- The less disruptive and less expensive it is to experiment, the more experiments there will be. Furthermore, more experimenters and the greatest breadth of knowledge increases the value of this experimentation. This is why architecture allowing user experimentation without altering the infrastructure of the network (such as end-to-end ideas) creates great value.
- Our theory is limited to situations where there are Business and Technical Advantages (BTA) that push designers and managers to favor central management and control. Examples are email (discussed in Section 6.2), centralized account management, which has advantages in terms of efficiency, tractability, and oversight, and centralized music distribution architectures that content owners prefer because of the better enforcement of copyright control.

This is an important assumption because the interesting case is when there are advantages to centralized structure, yet distributed architecture creates the most value. When the best of many experiments is greater than the business and technical advantages of central control and management, then flexible distributed architecture makes sense. This means market uncertainty is one important variable determining the value of network, protocol, and service/application infrastructure based on whether it allows distributed applications/services.

Below we present a model that quantifies this theory using techniques from extreme order statistics [7] and real options, which is based on work by Baldwin and Clark [6] about the value of modular design.

## 4.2 Model

Our mathematical model [2,5] validates the above theory from an analytical viewpoint allowing visualization of the tradeoff between market uncertainty **(MU)**, the value of experimentation from distributed structure, the advantages of centralized control and management, and how these networks, protocols, and services/applications evolve over time.

### *4.2.1 Value of Distributed Structure*

Our model is based on the cost and value differential between more efficient centralized management and more flexible distributed structure. We assume a distributed network infrastructure enables end users (or organizations) to experiment with distributed services and applications, and that centralized management and control implies that end users are prohibited from any type of experimentation in the context of offering new services.

In this model, the Business and Technical Advantage (BTA) of a centrally managed architecture, relative to a more distributed management style, is represented as a cost difference. Our model applies when it costs less to provide services with centralized management and control than with distributed architecture; thus, the BTA is positive.

The expected value of a particular management structure is the total revenue minus the total cost of the network infrastructure, protocol structure, or service/application architecture. With a central management structure that is restrictive to end user experimentation, we assume only one experiment instance (because it is so hard to experiment). Thus, the value is the expected total revenue (i.e. V the distribution mean) from a single experiment, minus the total cost.

With distributed structure that promotes end user experimentation, there are $n$ separate attempts to meet the market with a market selection process by users to determine the best outcome along with its market value. From above, we know Q($n$) denotes the value of parallel experimentation, with market uncertainty (i.e. the standard deviation) the scaling factor; therefore, the value of the best choice from n experiments with the benefit of parallel experimentation in uncertain markets factored in is the expected total revenue from the best of n individual attempts to meet an uncertain market minus the total cost of the winner.

The main difference between the value of centralized compared to distributed structure is the additional term representing the marginal value of picking the best of the $n$ experiments (i.e. MU*Q($n$)) along with the difference in cost between the architectures (i.e. BTA).

Figure 3 illustrates this result: it is the surface representing the value of users having many choices. It shows the value of picking the best experimentation (z-axis) along with its relationship to both market uncertainty (x-axis) and the number of experimental attempts (y-axis). When MU is low, a lot of experimentation is not useful, because the best experiment is only slightly better then the average (i.e., the left hand side of Figure 3). However, with high MU, the best of $n$ trials is likely to be significantly better than the average (the right side of the figure).

Experimentation by users is worthwhile if the expected value from the experimentation enabled by distributed architecture exceeds the expected value from centralized structure, which is equivalent to the benefit of experimentation with market selection being greater than the business and technical advantages to centralized control (i.e. **MU*Q(*n*) > BTA**).

The above argument illustrates a real options framework to compare the value of distributed structure allowing a diverse group of innovators to a more restrictive centralized architecture. It explains how distributed architecture creates more expected value even when it has business and technical disadvantages when compared to centralized control and management.

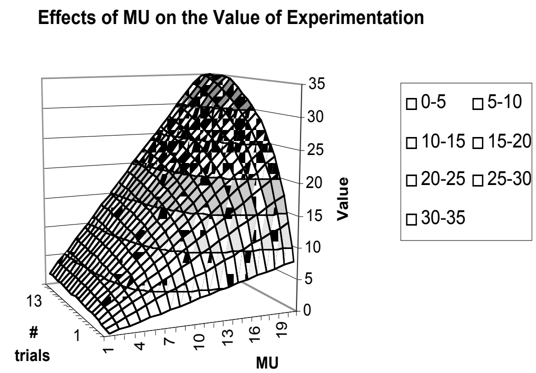


**Figure 3 Value of the Best of Many**

### *4.2.2 Learning*

The above model provides a framework to help understand and visualize the relationship between market uncertainties, many parallel experiments, and the advantages of central management. By modeling architectures that evolve from generation to generation, this section expands our basic model. This is accomplished by introducing learning–that is, experience from previous generations of experiments about the preferences in the market.

The effect of learning is to flatten the distribution curve by decreasing the standard deviation (i.e. market uncertainty). Learning reduces the benefit of many experiments because each experiment falls within an increasingly narrowing range centered around the mean (See Figure 2); thus, over time, many experiments help less and less. To incorporate learning, we introduce a model based on difference equations, which utilizes a learning function [2] to decrease market uncertainty at each generation. This multi-stage model captures the evolution of infrastructure over time.

One important question is whether it is better to have fewer generations with more experimentation per generation, or more generations 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 might be best. However, if MU does decrease, it limits the gain from experimentation, thereby making the answer dependent upon the rate of decrease.

## 4.3 Believability of Results

These results display a smooth, somewhat linear surface, which is attributed to good behavior of the normal distribution. This distribution is well suited to visualize the tradeoffs between efficiency, flexibility, evolution, and market uncertainty. It is unlikely that experiments in real life will follow a normal distribution and will be un-correlated as we assume. However, our model mostly depends on the best of many experiments always exceeding the expected value of a single experiment; we believe that for most realistic distributions this is true. We suspect a distribution with a heavy tail on both ends, which implies that with enough experimentation, there will be tremendous failures and stunning successes. The link between the ability to experiment and market uncertainty does not depend on using the normal distribution that the above results are based on.

## 5. REAL NUMBERS FOR BTA/MU

Below we discuss deriving real numbers for the business and technical advantages (BTA) of centralized control and the market uncertainty.

### 5.1 BTA

With careful cost accounting, the advantages of centralized management are not hard to ascertain. BTA is the total advantage achieved with central management. It includes both management and technical components. BTA is very general, as it must capture all the advantages of centralized architecture. The tangible advantages (such as less people to manage) are directly convertible into hard numbers. Less tangible advantages (such as the ability to monitor network use) can be converted into a cost advantage. Computing BTA is mostly a finance and accounting exercise.

### 5.2 Market Uncertainty

There are several techniques to determine the level of market uncertainty. We also propose several methods to put hard numbers to market uncertainty.

#### *5.2.1 Course Grain Estimate of MU*

Previous work has examined different metrics to gauge the level of market uncertainty: ability to forecast the market [8], emergence of a dominant design [9], agreement among industry experts [5], feature convergence [5], commodity nature of a product [5], and changes in standards activity [5].

#### *5.2.2 Fine Grained Estimates of MU*

How might real numbers be assigned to market uncertainty? Market uncertainty is the standard deviation of the distribution describing the value of services or applications for a particular purpose. There are two ideas we are working on in order to better quantify what market uncertainty means. One method is valuing services provided with open garden business models, such as iMode. Open garden structure promotes experimentation by outside (and independent) service providers and vendors. Another possibility is to look at the market capitalization of organizations that provide Internet services. Can this type of information help determine a number for market uncertainty? We think so!

## 6. EVIDENCE OF CORRECTNESS

If our theory/model is correct, we expect user preferences to change in relationship to the market uncertainty. When market uncertainty is high, or is changing from low to high, we expect users to prefer distributed solutions; when market uncertainty is low, or changing from high to low, we expect centralized architecture to become more popular. We illustrate this with two case studies [5]: voice services and email.

### 6.1 Voice Services

There are two ways for business to provide voice services: a distributed architecture using a PBX located on site, or a centralized structure using the Centrex service offered by traditional phone companies. Figure 4 is the number of installed Centrex lines from two different sources [5]. It illustrates the unpredicted resurgence of Centrex in the mid-80's. Gaynor's thesis argues that a decrease in market uncertainty was one important factor influencing the shift back to Centrex services at that time. History shows that successful PBX features such as caller-ID, voice-mail, and automatic call distribution have migrated from PBX to Centrex. About this time a general agreement had developed about what phone service consisted of, and whatever the method of delivering the service, the service itself was mostly the same. This example illustrates how successful ideas from distributed environments can migrate into more centralized structure, and in this case, more efficient structure when market uncertainty is reduced.

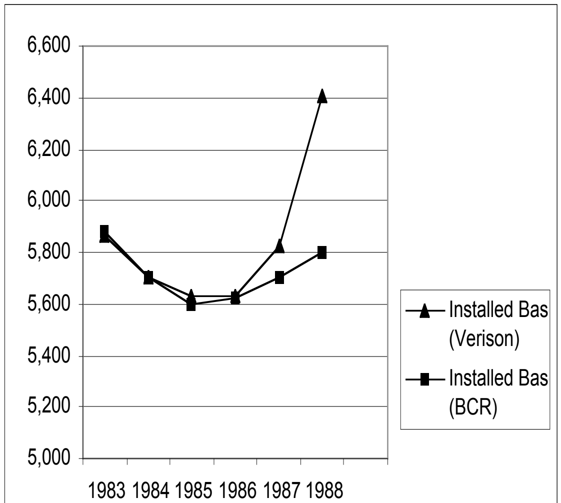


**Figure 4** Centrex Installed base (in 1000's) of lines

### 6.2 Web Based Email

Email has become one of the Internet's killer applications over the last ten years. Initially, email based on Internet standards had a distributed flavor with traditional POP based architecture; there has, however, been a recent shift in user preferences. As Figure 5 illustrates, centralized web-based email architecture boasts explosive growth. Gaynor's thesis argues that this shift to centralized web-based email is linked to a reduction of market uncertainty in the email area, as evidenced by the fact that email user agents began to offer the same set of features, the emergence of a dominant design (i.e. the Internet set of email standards) and stability in the base standards. Centralized email became popular after the uncertainty with email standards decreased.

Both the voice and email case studies support our theory: market uncertainty is a critical factor in deciding between centralized or distributed management structures.

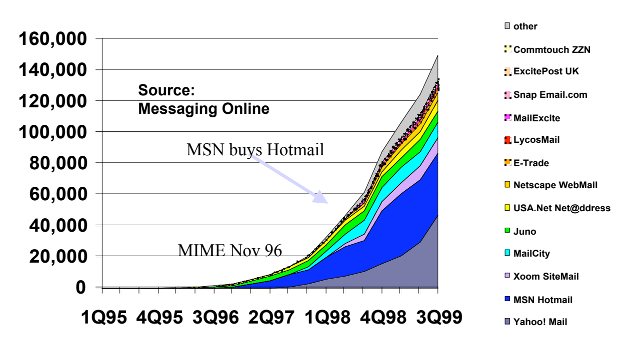


**Figure 5** Growth of web-based Email (in 1000's) of lines

# 7. APPLYING THIS FRAMEWORK

There are several ways to use this real options framework: as a visualization tool to help management understand that there are strategic advantages to a more costly, but more flexible, network infrastructure, as a quantitative guide for technical design choices[1], and as a tool for VCs to evaluate investment options. Venture Capitalists, Marketing and other non-technical managers need evidence that building more expensive architecture can translate into more revenue, since users are willing to pay more for services that meet their needs better. In the examples below, we apply our model as a lever to convince management that choices such as end-to-end structure, SIP as a VoIP protocol, and open gardens create tremendous value when markets are uncertain. Our framework is not yet ready as a quantitative guide to help technical designers, however, once we better understand how to estimate market uncertainty (See Section 6), this framework will be more useful in this context. In its current state, our model provides an intuitive justification of the value of distributed structure; in the future, this framework could become a valuable quantitative method to allow technical network engineers to evaluate different proposals for building infrastructure and applications.

The Internet has changed, which has caused a re-thinking about end-to-end ideas. VoIP, Broadband, and other new technologies are posing many important questions that need answers. Data and voice are converging, but the dominant standards have not emerged to achieve this merging of data types. Network owners have choices about supporting open or walled garden business models. The decisions of today affect the value of innovation tomorrow. Below is a brief look at these new technologies. We have discovered a common thread: network and application infrastructures that allow distributed end-to-end experimentation, as well as efficient centralized control of the most popular service/application, are doing well in the current market.

## 7.1 Today's Internet

The end-to-end agreement has been fundamental to the success of the Internet. Because of the historical end-to-end structure of the Internet, end users have been able to experiment with new ideas, unlike in the centrally controlled PSTN. User experimentation with end-to-end services/applications translates into more choices for all users. Our real options argument explains the value of allowing users to experiment, as well as how this value is linked to market uncertainty–high market uncertainty implies greater value of allowing end-to-end services/applications.

Today's Internet is different from its early days because of the scope of its use and those who use it. The end-to-end argument that fostered innovation in the early Internet must adapt to its current reality as Clark's and Blumenthal's [10] recent paper discusses. They explain the impact of new network devices such as Network Address Translators (NAT) [11] and firewalls that affect the end-to-end nature of Internet applications. NATs and firewalls, by limiting what users can do, adversely affect innovation within the current Internet. Other protocols, such as a service authenticating a user before allowing them to perform a task, are not end-to-end in the pure sense, yet can be seen as a combination of end-to-end transactions. Even more important, this type of verification from a "trusted authority" makes intuitive sense in today's Internet, and does not appear to limit experimentation. We find that some alterations from pure end-to-end services still retain the general ideas and still promote experimentation, but other alterations, such as NAT's, are not justified and break the end-to-end paradigm in unacceptable ways because they stifle innovation.

## 7.2 NATs and Firewalls

NAT architecture breaks unique global IP addressing. What might have seemed a clever idea was not, because NAT's breaks the end-to-end model when changing the network address within the IP header of a data packet. Using a NAT prevents a whole range of existing and future services/applications from working. Any service that requires knowing the IP address of one endpoint (as seen by another endpoint) does not work with a NAT[2]. For example, most Voice-over-IP protocols, including SIP and H.323, break with NAT technology because the IP address of each endpoint for these voice applications to work must be known by the other endpoint. NATs break end-to-end Internet security protocols such as IPSec, because part of the underpinning of IPSec is the uniqueness and security of IP addresses, and NATs change this. While not prohibiting innovation, NATs do make it more difficult.

In the context of innovation, however, firewalls have a greater negative impact than NATs. First, while NAT is not required in a firewall product, most have this NATing ability. More important, the basic function of a firewall is to filter out potentially bad traffic[3], which is often defined as anything to a transport level port address not on the list of allowed services. This is far worse than the NAT function, because instead of breaking a particular class of applications (like NATs do), it blocks all new applications not on the approved list. Innovating new applications with firewalls between users is not possible unless the firewall is configured to allow experimentation of this new service. The filtering function of firewalls is bad, but when combined with NAT, innovation becomes even more constrained.

While they may be attractive and are sometimes required for other reasons, firewalls and NATs make network

[1] Thanks to an anonymous reviewer for this comment.

[2] The IETF has just approved a protocol called STUN [15], which allows a device behind a NAT to know it's external IP address.

[3] Firewall filtering is not as secure as some believe. The technique of HTTP tunneling described in the Aprils fools' day RFC 3093 explains why.

experimentation harder, and thus, according to our real-options framework, reduce the overall value of the network since users will have fewer choices.

Figure 6 illustrates one aspect of the architectural evolution of the Internet: the degree of centralized control and management. It illustrates the traditional Internet and its current direction. The traditional (end-to-end) Internet had a very distributed infrastructure, which implies it does best when market uncertainty is high. Current architecture changes to the Internet such as MPLS, NATing, firewalls, Megaco, RSVP QoS services, and security services are pushing the current Internet infrastructure toward the right because of the imposed centralized control of these architectural changes.

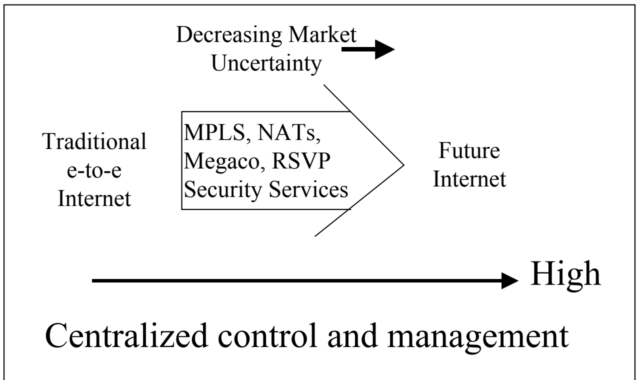


**Figure 6 Evolving Internet**

## 7.3 VoIP – Why SIP is Winning

It now appears that SIP will become the dominant standard for end device control in VoIP at least for the near future. Our real options based framework illustrates why architectures such as SIP, which allow the flexibility to build infrastructure with either end-to-end or centralized structure, have more value than standards such as Megaco/(H.248) that force centralized control by prohibiting end-to-end services [12], or proprietary protocols such as Cisco’s Skinny where a single vendor controls the standard. SIP promotes more innovation by enabling end users to experiment, through intelligent end devices, which provide services. This creates more choices for users, and according to our real options framework, creates more value at least until such a time when a common understanding will have developed on the nature and features of Internet-enabled voice services.

Figure 7 illustrates real options predictions of the value of different architectures for voice infrastructure. With traditional voice services, the market uncertainty is low. Centrex, a centralized architecture, allows very little experimentation; thus its value falls in the front left corner, the region of lowest value. With VoIP, the market uncertainty is high, which increases the value (moving to the right side). Proprietary protocols make innovation hard because experimentation is limited to a single vendor, which places its value in the right front. SIP is in the high value region because it promotes experimentation, and the high market uncertainty makes the experimentation worthwhile. Our frameworks prediction that SIP creates the most value seems to be correct as SIP is winning the standards battle for controlling end devices that provide VoIP services.

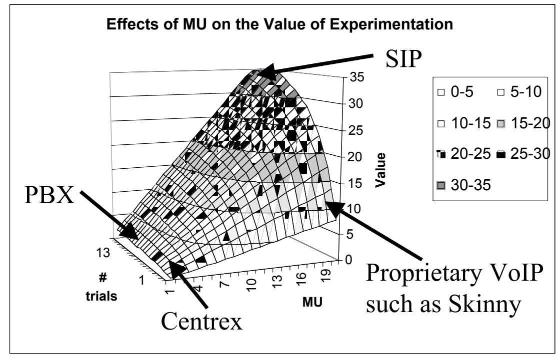


**Figure 7 Mapping VoIP architecture into a real options framework**

## 7.4 Open Vs Walled Garden

This real options framework is useful to evaluate policy decisions. Observers such as Larry Lessig [13] are applying end-to-end type thinking towards policies related to owners of cable networks and independent service providers that want to provide network services using the basic infrastructure of the cable network. Should cable network owners be able to control what services (and hence control, among other things, what content) a user can choose from? Our model implies that this is not a good idea even for the cable network owner. As we discussed earlier, greater value emerges when users are able to choose the services and the providers of these services without interference from cable network owners. Any other policy stifles innovation, causing a reduction in value. Professor Lessig agues how important it is to keep neutrality in the Internet. He notes that many companies such as Disney and Microsoft support this open garden view of the Internet. Today's telecom policy regarding open cable network access will sculpt tomorrow’s landscape; it is critical that policy makers choose the right policies. As the Internet has shown, “right” means allowing end users to choose the services and content they want.

Cable network owners argue that because they paid for building the network infrastructure they should have complete control over user access to all content and services. However, the real options framework illustrates how when users’ choices are limited by the network owner or operator, everybody loses. Achieving maximum value in the context of meeting user needs requires service providers unaffiliated with the core network to be allowed to offer content and services/applications

One measure of a networks value is the satisfaction of its existing users. Another measure is the network’s ability to attract new customers. Our real options framework illustrates why open gardens have greater value (measured by either metric) than walled gardens. More users with greater satisfaction will translate into increased revenue. Users like open gardens because there are more choices for them than in a walled garden that limits their choices. Users living in open gardens are less likely to switch networks to find service or content more to their liking than users in walled gardens because those in open gardens already have these choices. In the context of value to current users, ability to keep users, and attractiveness to new users, our model illustrates why open gardens win every time (unless you assume that there will be no further innovation of network-based services in the future).

## 7.5 Common Thread

The common thread in all these examples is the success of protocols and infrastructure that allow flexibility. The flexibility of distributed services promotes innovation, as well as centralized control, so that successful applications can be efficient, scalable, and safe. At many levels of the protocol stack, protocols and applications being adopted are flexible in regard to the architecture of applications/services they allow. At the application protocol layer, SIP and Internet email protocols demonstrate management flexibility. SIP allows true end-to-end services as well as a more centralized architecture via proxy servers and calling agents [12]. Internet email allows both distributed and centralized management structure as discussed in [5]. Finally, at the highest level are Web-based user applications, which have the property of flexible management structure. Protocols that are flexible in regard to management structure are successful at many layers, which highlights the value of flexibility in allowing services/applications to have either distributed structure or centralized control.

# 8. RELATED WORK ON OPTIONS

The framework used to illustrate why high market uncertainty enhances the value of allowing distributed services is based on the theory of real options, which extends the theory of financial options to value options on real (non-financial) assets [14]. Real options provide a structure linking strategic planning and financial strategy. Similar to financial options, real options limit the downside risk of a design (or investment) decision without capping the upside potential. This theory has proven useful in examining a plethora of situations in the real world, such as staged investment in IT infrastructure [14], oil field expansion, developing a drug [14], showing the value of modularity in designing computer systems [6], explaining the value of modularity in standards [3,4], and network based services [5]. This paper demonstrates how network infrastructure that allows distributed end-to-end services/applications is needed to maximize value in the context of meeting users needs in uncertain markets. Previous work on real options by Baldwin and Clark [6] discussing the value of modularity in a computer system provides the theoretical framework for this model, which is expanded upon by Gaynor and Bradner [4][5].

# 9. CONCLUSION

This paper presented a real options framework useful to compare network, protocol, and service/application structure. Our model explains why high market uncertainty implies that service providers and users will profit from a distributed architecture because of the value of experimentation, as well as the desire of users to have choices. It also illustrates how when market uncertainty is low, an efficient, safe, and controllable centralized management structure is more desirable to service providers and users because users don't value many choices when all the choices meet their needs well. We present a theory and model, which is supported by empirical evidence. Finally, we illustrated the value of this framework by applying it to several diverse areas such as the value of end-to-end ideas, analyzing the choices for VoIP architecture, and explaining the value of open compared to walled garden service/application business models that promote users choices in content/services they have access to. Our model explains why users prefer infrastructure that allows flexibility in usage.

# 10. REFERENCES


[1] Saltzer, J, and Reed D, and Clark, D., "End-To-End Arguments in System Design.," ACM Transactions in Computer Systems 2, 4 , Nov: 1984 p 277–-288.

[2] Gaynor, M, The Effect of Market Uncertainty on the Management Structure of Network-based Services. Ph.D. Thesis, Harvard University 2003.

[3] Gaynor, M., Bradner, S. Iansiti, M, and Kung, HT., The Real Options Approach to Standards for Building Network-based Services., Proc. 2nd IEEE conf on Standardization and Innovation, Boulder, CO,. Oct, 2001.

[4] Gaynor, M. and Bradner, S. Using Real Options to Value Modularity in Standards, Knowledge Technology and Policy, Special on IT Standardization, 14:2.*, 2001

[5] Gaynor, M, Network Service Investment Guide: Maximizing ROI in Uncertainty Markets, Wiley, 2003.

[6] Baldwin, C. and Clark, K. (1999). Design Rules: The Power of Modularity. Cambridge, MA: MIT Press.

[7] Lindgren, B.W, Statistical Theory, NY: Macmillan, 1968

[8] Tushman, M, and Anderson, P, "Technological Discontinuities and Organizational Environments," Administrative Science Quarterly, 31 (1986) pp 439-465.

[9] MacCormack, Alan, "Towards a Contingent Model of the New Product Development Process: A Comparative Empirical Study." Working Paper 00-77, 2000, HBS.

[10] Clark, D, and Blumenthal, M, "Rethinking the design of the Internet: The end-to-end arguments vs. the brave new world,". TRPC August 10, 2000.

[11] K. Egevang and P. Francis, RFC1631 – The IP Network Address Translator (NAT). May 1994

[12] Gaynor, M, "Linking Market Uncertainty to VoIP Service Architecture", IEEE, Internet Computing, July/Augest 2003

[13] Lawrence Lessig. 2002., The Future of Ideas: The Fate of the Commons in a Connected World, Vintage Books., 2002

[14] Amram, M. and Kulatilaka, N., Real Options, Managing Strategic Investment in an Uncertain World. Boston, MA: HBS press ,1999.

[15] Rosenberg, J, and Weinberger, H, .. RFC 3489, STUN, IETF, March 2003.