Using Real Options to Value Modularity in Standards

Mark Gaynor and Scott Bradner

Mark Gaynor received a Ph.D. in Computer Science in June 2001 from Harvard University and is now an assistant professor of Information Science at Boston University, School of Management. Gaynor's research has included high-speed ATM networks, using computers and networks as teaching aids, models of standardization, and economic models of network-based services. Before graduate school at Harvard, Mark was a software engineer, network designer, and manager at Vitel, a small international telecommunication company located in Boulder, Colorado. While at Vitel he received a ME in telecommunications from the University of Colorado at Boulder. He may be reached at <mgaynor@bu.edu>.

Scott Bradner is the co-director of the Transport and Sub-IP areas in the IETF; he is a member of the IESG, and until June 1999, was an elected trustee of the Internet Society, where he still serves as the Vice President for Standards. He was also co-director of the IETF IP next-generation effort and is co-editor of "IPng: Internet Protocol Next Generation" from Addison-Wesley. He is a member of the Wiley Network Council. Mr. Bradner is a senior technical consultant at Harvard University, Office of the Provost, where he provides technical advice and guidance on issues relating to the Harvard data networks and new technologies. He founded the Harvard Network Device Test Lab, is a frequent speaker at technical conferences, a weekly columnist for *Network World*, and an independent consultant.

This paper proposes a model of technology standardization based on modular standards and the effect of market uncertainty on the value of modularity in standards. A real options model quantifies the value of modularity in standards, illustrating that a rational way to standardize some IT technology in uncertain markets is with modular design, end-2-end structure, and proper staging of the standard. Highly modularized standards provide a higher option value because of the ability to select the best modules to change at a fine granularity.

Introduction

This paper is an expanded version of our previous paper (Gaynor and Bradner 2001). Its main contribution is a model illustrating the value of experimentation with modules of a standard that have the most potential to affect the total value of the standard. Using a model based on real options, it shows the value of users having choices when market uncertainty exists.

Product development involving computers, software, and networking has had a profound impact on theories of innovation and product development. Technology in these products changes very rapidly when compared to that in many traditional industries such as automobiles (Clark 1985), or production of television picture tubes (Utterback and Suarez 1993). Just imagine autos that double in speed every 18 months, similar to the performance increase in microprocessors. The faster evolution of computer and information technologies do not fit traditional product development theories, which depend upon periods of disruptive innovation followed by less drastic incremental changes (Dosi 1982, Nelson and Winter 1977, Anderson and Tushman 1990, Abernathy and Clark 1985), since the period these technologies are stable is very short. A new breed of models (Brown and Eisenhardt 1997, Gould 1992, Iansiti 1995) views the evolution of technology as a continuum of changes, rather than as the punctuated equilibrium of the past.

Customer expectations co-evolve with technological change at today's faster pace, creating uncertainty in consumer preferences. Clark (Clark 1985) points out that when a new technology is born, customers have no education about the technology and tend to view it in the context of what it replaces. The evolution of customer expectations of the web is a good example. At first, the web was mainly a tool for researchers sharing information, the important service attribute being existence of the data and its accessibility by heterogeneous computer systems. Only later, as the interactive nature of the web matured, did consumers become more sophisticated in the services they demanded. Now information layout, e-commerce, and usability have become important attributes of web-based services.

Changes in product development in the computer age are parallel to alterations in effective standardization of technology. With slower-moving technologies, standardization occurs in the relatively stable period after technology selection. However, this stable period is short or non-existent in fast-changing technologies such as DRAMs, where useful standards must be timely, and produced in a few months instead of years (Rhoden 1999). The uncertainty created by evolving customer preferences means that (once created) standards must have the ability to evolve along with the end users of the standard.

Standards have become more important to business, thus causing strategic management of companies' standards policies to play an increasingly important role in formulation of overall corporate strategy. Evidence of this is the increased memberships of fee-based industry consortiums and alliances such as World Wide Web Consortium (W3C), Asynchronous Transfer Mode (ATM) forum, and X/Open. A report to the chairman of W3C in 1999 shows an increase from 30 to 370 members. Cargill (Cargill 1997) also notes that at one point in the 1990's, consortiums for business-based standards increased by two per month. This shows that the demand for standards outstrips the supply produced by traditional Standard Development Organizations (SDOs) such as ISO and ANSI.

Networking, IT, and other technology standards in areas of uncertain user preferences are not static documents, but dynamic complex adaptive systems (CAS) that must interact and change within their environment. Factors causing these standards to behave as CAS are an increased number of users with diverse, fast-changing, unpredictable requirements, and the uncertainty of implementing a standard to work correctly by being interoperable with other independent implementations of the standard. For better success in this new uncertain environment, standards must follow principles similar to the way evolution and natural selection pick the fittest organisms, but, with the market[1] as the selector picking the fittest technology in terms of the users (or users and vendors). This paradigm is well suited to describe the standardization process in today's ever-changing dynamic environment.

A standard development methodology that promotes a broad range of experimentation combined with market selection will better serve end users by involving them in the standardization process. In high market uncertainty, promoting experimentation with new proposed standards and standardizing the technology adopted by most of the community decreases the risk of an unaccepted standard. Design principles, such as the end-2-end argument, that push intelligence to the network's edge, help promote this type of experimentation because they broaden the range of participants able to innovate by allowing users to create services.

Thinking in terms of options (See Section 2.2) means considering how uncertainty affects the value of flexibility. It shows how delaying decisions until future information becomes available increases the expected value of an investment. High uncertainty increases the value of this management flexibility. Applying options thinking to standards gives the right, but not the obligation, to follow a path of standardization. The path followed by a standard depends on how the market unfolds. For example, the IETF category "proposed standard" is a profile of options; it allows the Internet community to choose the standards that succeed by exercising the option to implement the standard and provide the services enabled by the standard.

In this paper, we propose a prescriptive model of technology standardization under uncertainty and show how its value is quantifiable using the theory of real options (Amran and Kulatilaka, 1999), a proven methodology for management of non-financial assets under uncertainty. Our model is simple and intuitive: start with simple standards structured in a modular layered architecture, then let the standards evolve, with the market acting as the selection mechanism. Our model of standardization shows how modularity creates value, and how market uncertainty increases this value. We explain how to apply this framework to the development of communication protocol standards, but do not provide a numerical example. Our model shows that modularity (up to a point) intrinsically creates greater value than an interconnected design. We argue that an end-2-end structure (see Section 4), where the network only provides basic transport services by pushing applications to intelligent endpoints, creates an environment conducive to experimentation. For example, the Network layer in a network protocol stack (IP) should contain the fewest protocols that provide only the most basic services, while the application layer should contain the most protocols to offer the most diversity in terms of services offered. Next, we discuss the value created by applying the methodology of introducing protocol suites (and protocols) with a few simple modules, and evolving the standards by creating new protocol modules or altering existing ones. Our theory shows that the evolutionary approach to development of entire protocol stacks, and protocols within each layer, maximizes the expected value of the standard.

This paper should be of interest to both academics, and practitioners interested in standards or the architecture of protocols. For the academic, this paper presents a new idea: using the theory of real options to value modularity in a standard. This paper is the tip of the iceberg; there is much further research in this area. For those who create standards, this paper presents a new mindset - think in terms of keeping the options open for further standardization; the more flexibility allowed, the better the expected outcome. We do not mean standards should be option-filled, but that the structure of standards should allow evolution of the standard as the environment changes.

The structure of this paper is as follows. Section 2 is the methodology showing how this theory is based on previous work about the way markets select standards. It discuses how the theory of real options has been used to show the value of modularity in computer systems design, when the technology outcome of system components is uncertain. In Section 3, we discuss uncertainty in the context of standards. Section 4 discusses the end-2-end argument and why it promotes innovation. Next, in Section 5, our model explains how the theory of real options helps quantify the value of modularity. Section 6 examines the evolution of modular standards in conditions of fast-changing complex technology. Section 7 contains empirical evidence of our arguments. Last, in Section 8, general rules of protocol standardization are discussed, and generalized to other standards.

Methodology

This work is theoretical and draws from two main areas of research. First, work by Vercoulen (Vercoulen and Van Wegberg 2000) discusses modularity in standards, and its effect on the selection of standards in dynamic complex industries. Next, research related to the theory of options shows the value of modularity and choice between the modules in computer systems. While based on theory, the empirical evidence supports our model (at a high level of analysis - see Section 7). The successful Internet protocol suite has a modular structure promoting end-2-end applications, they started simply, and then evolved in complexity as user needs became more stable.

Modularity and Selection of Standards

Vercoulen (Vercoulen and Van Wegberg 2000) discusses how modularity in standards adds value by creating standards that are complex, but able to react to dynamic change quickly. It discusses how modular standards may work best with a combination of market and negotiated selection. Negotiations sometimes help develop complex standards that fit together, but this negotiation process may be too lengthy for dynamic markets. By combining both selection modes, complex working standards can be created in a timely manner. This work classifies complex modular standards as complex dynamic systems and builds a base for our theory by illustrating that even complex technologies can be standardized with modular standards such that the complexity of any module is low relative to the standard as a whole.

Modularity of complex standards has advantages and disadvantages. Vercoulen points out benefits of modularity such as: specialization where different parties develop different modules, scalability of the system, and innovation by including new modules. Also included are the negative aspects of modular systems such as: coordination failures are possible between modules, resources required to link and coordinate modules add to system overhead, and connecting modules into a cost effective system is non-trivial.

A main way to accomplish modularity in standards is by specifying the interface between the modules of the standard. This defines the data and its format that passes between modules. For example, in the Internet protocol IP, the standard specifies what a transport protocol such as TCP must pass to IP, and what IP provides to the Link layer below. Our work focuses on the advantage modularity gives to innovation, while accounting for the additional expense of the modularity as discussed in Section 0.

Theory of Options

The theory of options has proven useful for managing financial risk in uncertain environments. To see how options can limit risk, consider the classic call option. It gives the right, but not the obligation, to buy a security at a fixed date in the future, with the price determined at the time the option is purchased. Buying a call option is the equivalent of betting that the underlying security will rise in value more than the price of acquiring the option. The option limits the downside risk, but not the upside gain, thus providing a non-linear payback, unlike owning the security. This implies that options provide increasing value as the uncertainty of the investment grows (i.e. as variance in the distribution describing the value of the security increases), since the downside risk is capped without limiting the upside potential.

Figure 1 shows graphically how this works. The non-linear payback of the option is the solid line, while the linear pay out of owning the stock is the dashed line. The option holder is able to look at the price of the security when the option is due and decide whether to exercise the option to buy the stock. It is the historical variability of the stock price, not the security price that determines the value of the option. This protects the option holder by limiting the loss to the cost of acquiring the option, no matter how low the stock price falls. Some risk-averse investors prefer this type of non-linear payback that caps the downside risk, but leaves the upside gain unaltered.

This theory of options is extendable to options on real (non-financial) assets (Amram and Kulatilaka 1999). Real options provide a structure linking strategic planning and financial strategy. Similar to financial options, real options limit the downside risk of an investment decision without limiting the upside potential. In many cases, this approach shows a greater potential expected value than the standard discounted cash flow analysis performed in most corporate environments. This is because the discounted cash flow methodology depends on computing the present value of a series of future cash flows. These flows must be determined at the time of the analysis. This makes it impossible to factor in the effect of flexibility in management that allows it to adapt to changing conditions. This theory is useful in examining a plethora of situations in the real world such as staged investment in IT infrastructure (Kulatilaka, Balasubramanian and Storck 1999), oil field expansion, developing a drug (Amram and Kulatilaka 1999), and even showing the value of modularity in designing computer systems (Baldwin and Clark 1999).

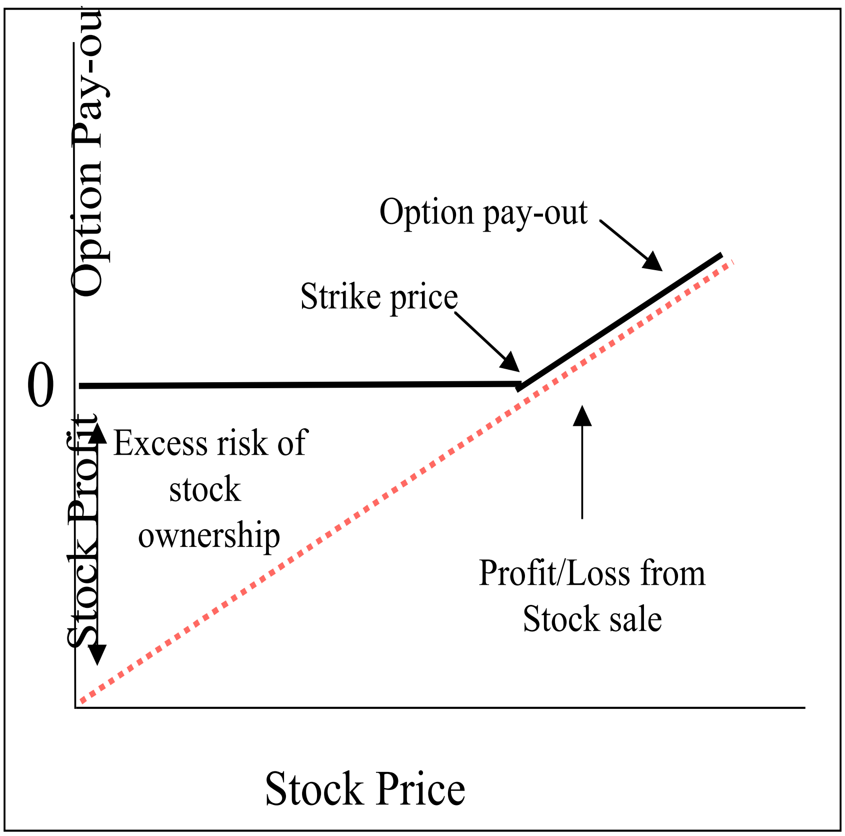

Figure 1 Option Pay-out

Staging the investment required to build large IT/Telecommunications systems provides an option at each stage of the investment. The option is whether to continue the investment or not, and is based on the most current information available about the uncertain market, economy, and technical attributes of the project. Starting out small, and evolving the project at various stages allows making more focused and relevant decisions, which in turn increases the expected value of a staged implementation over that of the single stage scheme.

In "Design Rules", Baldwin and Clark apply this theory to study modularization in the computer industry. They show how modularization of computer systems design (like the IBM 360) has changed the industry tremendously. A modularly designed computer consists of components that have defined interfaces. Because each component conforms to its interface rules, modules that follow the defined interface are interchangeable. In contrast, an interconnected system has no swappable components because only a single massive component exists.

To see how a modular design provides value, consider the evolution of a typical computer system. When redesigning a computer that has its functional pieces interconnected, the new artifact provides a single choice; the new system performs as a whole either better, or worse than its predecessor does. However, with the modularized version, the designer has the option to include each new module created for the next version, on a module-by-module basis.

This modularization allows the designers to experiment on modules that have the most potential for altering the value of the system. Each experiment is an alternate design of the module. Performing many experiments on the components most critical to overall system performance maximizes the overall value. Because of the modular design, the designer now has the *option* to pick the best outcome from many trials. For example, suppose the designers of a new computer system need to increase the rate a CPU module processes instructions. By attempting several technically risky new technologies for a CPU, the designer can improve the odds of reaching the goal of faster instruction execution. Futhermore, if all the risky designs fail, the old CPU is still usable because the design is modular. The modular design allows using the old CPU, but also the *option* to include any improved components such as the display or memory systems. This approach is impossible with the interconnected version: the only *option* is to take or leave the entire new system. The modular design increases value by providing a portfolio of options rather than the less valuable option on a portfolio (Merton 1992).

Other authors have used option theory to address a similar problem of supporting information infrastructure decisions in market uncertainty. Taudes, Feurstein and Mild (Taudes, Feurstein and Mild 1999, Taudes, Feurstein and Mild 2000) examine the evolution of software platforms when the future applications are unknown. They show how an options pricing model overcomes the limitations of using the discounted cash flow/ net present value methodology because of the uncertainty in the market.

Uncertainty with Standards

Modern technology is complex, with much uncertainty; users and vendors each have different needs for the standardization process and competition exists both between and within standards. The regulatory environment is now different, allowing support for a more open standardization process, and providing incentives for those creating the standards. Our theory depends on competing technical solutions for standards to provide the users with options. Unpredictable and dynamic user needs cause vendors to have incomplete knowledge of how a particular standard will mature and be used. While general services are predictable in many instances, the particular architecture, feature set and implementation that is widely adopted is often not predictable. Email is a good example of this; the demand was clear, but it took several generations of competing service offerings to converge to the Internet standards-based solution. As noted by Clark (Clark 1985), customers do not have the education with new technologies to understand the possibilities. These customer expectations must evolve along with the technology; the interaction between the technology and consumer preferences is very complex. Similar to new views of product development (Iansiti 1995, Gould 1992), standards fit into the context of

Complex Adaptive Systems (CAS) (Holland 1992, Kauffman 1993, and Waldrop 1992). This implies effective standards must evolve and have a selection process to pick from many competing options.

Uncertainty about the success of a proposed standard is one reason that standards need to start out simple, but display the flexibility (the cost of this flexibility is discussed in Section 0) to evolve within a continuously changing environment. For our model, we only examine uncertainty (and the associated risk) in market prediction and do not address economic or technologic uncertainty.

The market for "standards-based products" can be very dynamic and hard to predict. We are interested in need-driven standards where users' needs are a moving target, making even accurate short-term predictions difficult. It may seem a contradiction to claim that demand-pull standards can have elements of unpredictability, but consider email, which was clearly a predicted success. However, the standard that became popular was not the predicted X.400 suite, but the Internet scheme. Many services have clear demand, but there still exists uncertainty about the particular feature set and implementation that will work best for most users. Firms sometimes get it wrong even in a well-defined market.

There are many examples of how vendors are unable to predict what will happen in today's world. Nobody guessed the WWW would be the "killer application" that popularized the Internet, or the dramatic impact the WWW is having on society. The success of the entire Internet and the value created by it vastly surpasses any estimates its creators could imagine (even in their wildest dreams and post-Internet bubble). Technologies like ISDN, SMDS and ATM did not meet the predictions of the experts. OSI transport, the suite of communications protocols developed by the ISO and championed by all the major vendors and governments (including ours), is dead. These examples show the complexity of the standardization environment for networking systems. Thus far, it has been hard to predict which standards will become successful and which ones will fail.

Even successful standards mature in unforeseen ways. Frame-relay is successful in low bandwidth application, but was developed as a medium- and high-speed WAN service. ATM failed to reach the desktop (in terms of ATM cells reaching the PC) as expected, but has become a viable solution only within the core fabric of high-speed IP routers, and recently in providing multiplexing in DSL links. It is precisely this unpredictability of which standards succeed, and which applications use what standards, that requires a new paradigm.

Linking Market Uncertainty to the Value of Experimentation

When vendors or service providers do not understand what users want, they must experiment with applications that have different feature sets. Each experiment is a product for the user and is one attempt to meet an uncertain market by anticipating the needs of the user. The economic value of experimentation links to market uncertainty by definition of market uncertainty - uncertainty is the inability of the experimenter to predict the value of the experiment. When uncertainty is zero, the outcome of any experiment is known with perfect accuracy. As uncertainty increases, the predictability of the success of any experiment's outcome is lower, because outcomes are more widely distributed. This means that successful products may generate large profits, while unsuccessful products may not. This link between experimentation and uncertainty is intuitive, as long as the definition of uncertainty is consistent with the variance of results from a set of experiments.

When market uncertainty is low or zero, the experimenter has a good idea of the market. This means that each experiment will match it well, and meet the needs of users. However, if uncertainty is large, then the experimenter cannot predict how the market will value the experiment. It may be a wild success (such as the Web), or a dismal failure, such as the attempt of PBX vendors to capture the business data LAN market in the 80's. Figure 2 shows how 100 experiments might be distributed on three examples of a normal distribution (each with a difference variance). The data points were simulated using an algorithm given in (Hiller 1967) for a normal distribution with mean = 0, and variance = 1, 5, and 10.

This shows how the value of uncertainty changes the benefit of experimentation; when uncertainty is low (variance = 1), as expected, the best experiment has a value around 2.5 away from the mean. However, when uncertainty is high (variance = 10), similar data yields a value of 25 away from the mean, an order of magnitude better than the low variance case. This shows how, in an absolute sense, as uncertainty increases, so does the possibility of performing an experiment that is a superior match to the market, as indicated by a value far above the mean of the distribution. The example shows that when uncertainty is low, even the best experiment is not far from the mean, but when uncertainty is high, the most successful experiment greatly exceeds the mean. The important point is that the difference between the mean of the distribution and the best experimental result grows as the standard deviation increases.

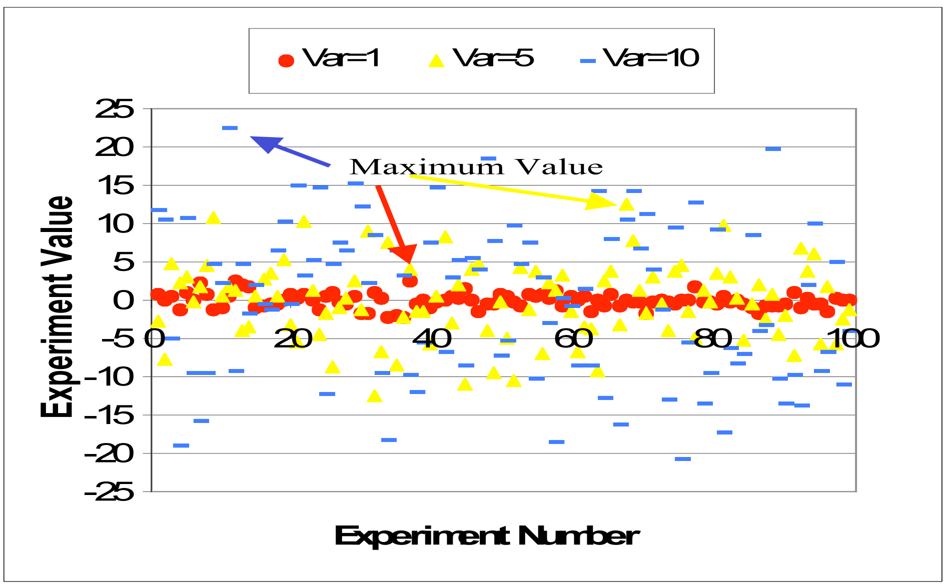

Figure 2 How Uncertainty Changes the Value of Experimentation

Figure 2 illustrates how these experiments are similar to a random walk, because each experiment is without direction. It assumes the experimenters have no focus, which is unlikely. This does not hurt our argument. We model the random walk approach as the worst case because it allows an equal probability of either a very bad result, or a very good one. When experimenters have more direction, we expect the distribution to shift to the right, improving the outcome.

How to measure MU

Precise measurements of market uncertainty (MU) may not be possible, but it is possible to estimate MU in terms of it being low, medium, or high. While difficult, estimating market uncertainty is important for showing the value of modularity. Fortunately, previous research by Tushman (Tushman and Anderson 1986) explores how to measure MU in terms of forecasting error, which is the ability of industry analysts to predict the industry outcomes. MacCormack (MacCormack 2000, MacCormack and Verganti 2001) discusses the existence of a dominant design as evidence of lower market uncertainty. These previous papers show the validity of estimating MU in the research community. For a more detailed discussion of how to measure market uncertainty, see the author’s thesis (Gaynor 2001).

End-2-end Structure

One idea that has helped the success of the Internet is its early promotion of end-2-end applications. Services with end-2-end architecture (Saltzer, Reed and Clark 1984) by definition have a distributed structure because they push complexity to the endpoints of the network. The idea is to keep the network simple, and build any needed complexity into the end, or edges of the network. Applications that are end-2-end are unknown to the network infrastructure. This means that changes to the network or permission to add new end-2-end services are not necessary, because nothing within the network knows about the new service. The end-2-end argument is one of increased innovation, and the proof of its validity is the success of the Internet with regard to innovation.

One main idea behind the end-2-end argument is that services offered by the network infrastructure should be as simple as possible. If you try to anticipate what services applications need, you will be wrong, and most likely inhibit new applications by constraining them to services that do not match their needs. Networks that only provide simple, basic services allow applications more flexibility in what they can do. The IP protocol in the Internet is a good example of this philosophy; it is simple, only offering the most basic type of network service - the unreliable datagram service. This simple core protocol has allowed immense innovation at the transport and application layers. Different applications modules can utilize the different transport protocols that match their needs, but all of them are built over IP, which has become the glue holding the Internet together. The success of the Internet is partially due to the simplicity of IP. Again, this validates the end-2-end argument.

By pushing applications to the user level with end-2-end applications, more experimentation is likely. There are several reasons for this. First, application-layer development is faster and less expensive than kernel work. Next, the pool of talent with the skills to do application-layer coding is greater. Finally, the participants allowed to develop new services are much broader at the application level because users can innovate, and as Hippel (Hippel 1998) shows, users sometimes are best suited to solve their own problems.

Since end-2-end applications do not require network infrastructure change or permission to experiment, users can and do innovate new services. Consider the creation of the Web. Tim Berners-Lee (Berners-Lee 1999) was not a network researcher searching for innovate ways to utilize the Internet, he was an administrator trying to better serve his users. He developed the Web to allow the scientists in his organization to share information across diverse computers and networks. It just so happened that his solution, the Web, met many other user needs. One powerful attribute of the end-2-end argument is that you never know who will think of the next great idea, and with end-2-end services, it can be anybody.

One reason the end-2-end argument is relevant today is that new services such as Voice-over-IP have a choice of management structures between an end-2-end architecture and a more centralized model. Session Initiation Protocol (SIP) is one way to provide Voice-over-IP. SIP can provide true end-2-end service, with nothing except the endpoints knowing that a conversation is taking place. Implementations of SIP work, but it is not the favorite model of traditional phone companies. The protocol known as H.248 (also known as megaco) is in co-development by the IETF and ITU-T. Based on the current architecture of the telephone network, it is a more traditional way to provide Voice-over-IP. It relies on a centralized server that coordinates the voice connections. This centralized server has knowledge of all calls it processes -- a nice feature for billing. It is not surprising that traditional phone companies support the H.248 model; they are sure it is superior to SIP. The disagreement about how to provide Voice-over-IP illustrates the importance of the end-2-end argument today.

A Model Illustrating the Value of Modularity

The first step towards quantifying the value of a protocol standard is to define what to value, and what metric to use. One value might be how easy it is to augment a protocol suite with new protocols standards and change existing ones. Another might be the value of a particular protocol in the context of how useful the protocol is for building services above it, or its scalability. Efficiency or speed of the implementation might be another good metric. These are not the only measures, but they are good examples. Our theory does not depend on the metric, but only on estimating the variance of the distribution describing its value. Real option theory provides a methodology to compute the expected values of giving users choice in uncertain markets.

Modularity

Our approach to value modularity in standards is similar to that used by Baldwin and Clark (Baldwin and Clark 1999) showing the value of modular design over its interconnected cousin in computer design. The advantage of using modularity within each layer of a protocol stack is similar to the benefits gained by using modular design in computer systems. It allows keeping the best new module for a protocol (maybe, picking the best outcome from many experiments) or keeping the old module, thus guaranteeing a higher expected value. To gain the most benefit from modularity there should be many choices for new modules. Architectures such as the end-2-end principal help provide many choices, because of the ease of experimentation and large base of users able to create new choices. We do not provide a numerical example of this theory, but refer the reader to *Design Rules* (Baldwin and Clark 1999) to see how to accomplish this.

Value of Modularity

Baldwin (Baldwin and Clark 1999) computes the value of modularity. Let V_1 be the value of a complex system built as a single module, and let V_j be the value of the same system with j modules. If we ignore the cost of modularity, then the value of dividing a complex system into j components: $V_j = j^{1/2}V_1$. That is, the modularized system exceeds the value of the interconnected design by the square root of the number of modules. This value does not depend on the variance of the distribution, because for each
module … for the new
module. … ed value is not
depende… variance of
the distr…

E…

A… expect to
happen l… bability of
experim… e of a
particul… xperiments
will fall … ard deviations,
and 2% … sults in Figure
2. To fir… ndard to find a
module … that finding
great sta…

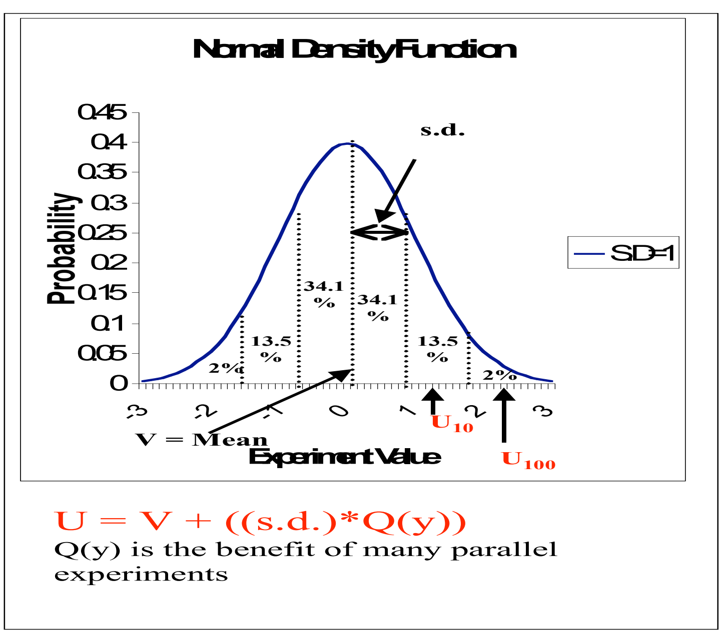

Figure 3 Best of Many Experiments

This figure shows U(10) and U(100), the expected maximum of 10/100 experiments. That is, U(10) is the value of the best experiment from a sample of 10 experiments. This maximum is composed of two different components: first is the effect of the mean, next is the offset from the mean. This offset from the mean (V) is itself composed of two parts: first, the effect of the standard deviation, and second, the effect of the parallel experimentation. Thus, we can express U(n) in terms of these parts: U(n) = V + Q(n)*S.D. That is, the maximum of n experiments equals the distribution mean plus the value of n experiments times the standard deviation of the normal distribution. Q(n) measures how many standard deviations from the mean U(n) is. Intuitively it makes sense that U(n) >= V, since to get an expected value of mean, we do the n experiments, take the first one (expected value = V), and disregard the rest. It also follows that the probability of U(n) greatly exceeding V increases as n or the variance grows.

Roughly for n = 2, Q(n) = .85, for n = 10, Q(n) = 1.5, for n = 100, Q(n) = 2.5, and for n = 1000, Q(n) = 3, again matching what is observed in Figure 2. The intuition behind this is that as you increase the number of experiments, the best of these experiments has a value that grows further from the mean, but at a decreasing rate.

As uncertainty increases, so does the gain from experimentation and thus the potential of profit. To see how this works, consider the following example: let the S.D. = 1, and n = 10 with a mean of zero. With n = 10, Q(10) = 1.5, so U = 1 * 1.5 = 1.5. However, if we increase the standard deviation to 2, then U = 2 * 1.5 = 3. This example shows that Q(n) is a measure of the number of standard deviations U is away from the mean.

This model, based on the best of many experiments, is option-based since a designer provides several options for a particular module. When only a single choice for a module exists, the expected value is lower than if the designer has many choices for the module. The model illustrates how uncertainty increases the benefit of many choices.

We assume a normal distribution and experiments that are not correlated with each other, but the basic idea holds for any distribution or correlation between experiments. The expected value of the best of many experiments may greatly exceed the mean, and is always at least as good as the expected value of a single experiment. The next section uses this result to model the expected value of modular standards.

A Simple Model of Modularity in Standards

In this section, we present one possible mathematical model based on our theory about the value of modularity and the extreme order statistics discussed above. This model shows the value of having many choices for any particular module in a standard, and how this value increases as the uncertainty of the effect of that module on the total system grows.

The uncertainty of a module in a standard is the variance of the distribution representing its value. MU is defined as the amount the module can influence the value of the total standard. As Figure 3 shows, V is the expected value of X, the random variable denoting the value of a module; that is, 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 module experiment to its designer is equal to the market uncertainty. This is because MU is defined as the inability to predict the value of a new protocol module to its designer.

As illustrated above, the marginal value of having more than one choice for a particular module is given by:

Equation 1 : U(n) = MU*Q(n)

Figure 4 is a surface that represents the additional value of having n choices for a particular module. It shows the value of experimentation along with its relationship to both uncertainty and the number of experimental attempts for a particular module.

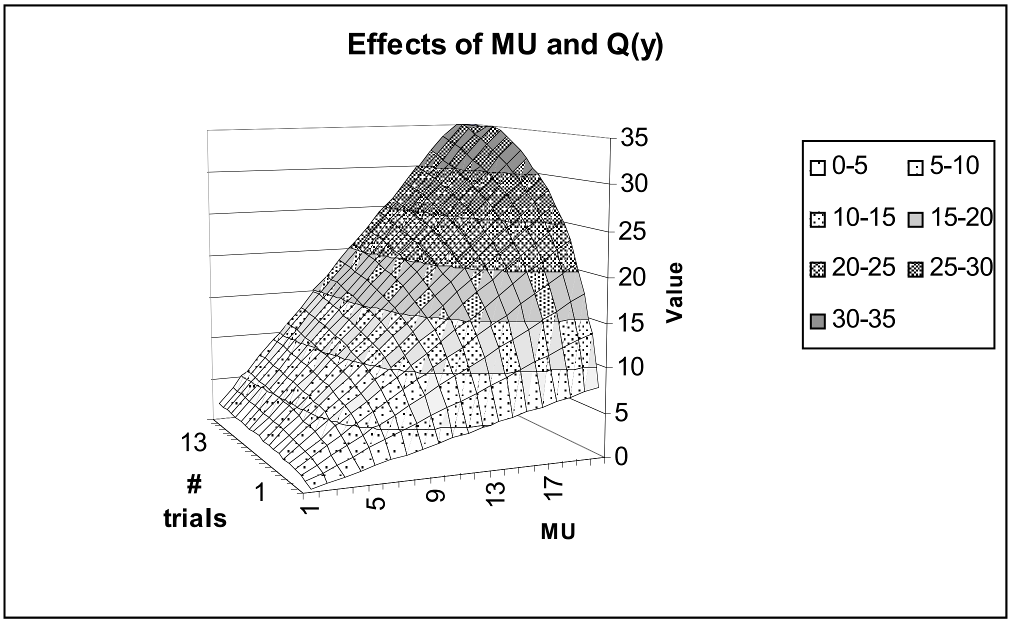

Figure 4 Value of the Best of Many

As expected, the value of a module to the total system performance increases at a decreasing rate with respect to n, the number of experiments. It is increasing at a constant rate with respect to MU, the uncertainty. The surface in Figure 4 illustrates this by showing the value (Z-axis) of running n (Y-axis) experiments with regard to the MU (X-axis). The curved lines for increasing n show the decreasing rate of increase, while the straight lines for increasing MU show the linear increase with regard to MU.

Value of End-2-end Structure

The above methodology makes possible the valuation of standards allowing end-2-end applications compared to standards that do not allow users to create applications. One example of two standards designed to provide a similar service, where one has an end-2-end architecture, are SIP and H.248, discussed in Section 4. SIP allows end users, service providers, and local developers to create new applications; H.248 only allows the service provider controlling the media controller and gateway to provide new applications.

In this model, we assume that service providers adopting standards promoting network infrastructure that runs counter to the end-2-end model must have some business or technical reason to do so. Reasons for this include management efficiency, protection of users from rouge applications, and the ability to track users. These and other reasons are discussed in (Gaynor 2001). This Business and Technical Advantage (BTA) of architectures that do not allow users to experiment, relative to the end-2-end model where users can create applications, is represented as a cost difference. BTA is the total advantage achieved by offering a service with standards architecture that only allows the provider of the basic network services to experiment with new applications. It 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 standards architecture L. E is for end-2-end type services, C stands for architectures that only allow the provider of basic network services to experiment.

This cost is comprehensive and includes both the internal and external components, including internal infrastructure, equipment (including software), and management.

For the services we are interested in : CP(E) > CP(C). It is more expensive to provide end-2-end services than to provide less flexible architectures that do not allow users to experiment. We are interested in cases where it make economic sense to allow end-2-end services even when it costs more to provide these services. When providing end-2-end services is less expensive our model is unnecessary. Thus, the equation for BTA is:

Equation 2: BTA = CP(E) - CP(C)

VP(L) is the expected value of a service with particular standards architecture L. This value is the total value the service provider receives for providing the service, minus the total cost of providing the service. For a service where only the network owner or manager can provide services, we assume they only make one attempt of providing the service. This means that there is only one choice for the module representing the new application. This implies the value of this application to its provider is:

Equation 3: VP(C) = V - CP(C)

For standards allowing end-2-end applications, we assume n application instances in a group of applications. Each application is a module that represents one product from a user or independent service provider. Then we allow market selection to pick the best outcome. Recall from Section 5.3.1.1, Q(n) denotes the value of parallel experimentation. Thus the value of the best service provided with the benefit of experimentation in uncertain markets factored in is:

Equation 4: VP(E) = V - CP(E) + MU*Q(n)

Allowing users to provide network applications is worthwhile if **VP(E) - VP(C) > 0** => **MU*Q(n) > CP(E) - CP(C),** which is equivalent to **MU*Q(n) > BTA**. This means it is better to allow users to provide services if:

Equation 5: MU*Q(n) > BTA

As market uncertainty increases, allowing users to provide services becomes more attractive because of the enhanced value of experimentation[2]. If the cost differential between end-2-end and non-end-2-end is less than the benefit gained from high market uncertainty and parallel experimentation, then the value of the best application from all the users and service providers is likely to greatly exceed the value of a single attempt to provide the service by the service provider that is also providing the basic network services.

Assuming that market uncertainty exists, Figure 5 shows the relationship between MU (the market uncertainty), BTA of using a central model (the business and technical advantage transformed into a cost differential) and n, the number of experiments run in parallel. This surface shows the relationship for a range of n (# of simultaneous service experiments) between 1 and 20. Points on the surface show where market uncertainty equals BTA/Q(n); the points above the surface show where end-2-end architectures work well because of the advantage of allowing many users to perform parallel experiments combined with market uncertainty. Points below the surface have low enough market uncertainty relative to BTA that the network manager is able to meet market needs with a single attempt. The forward edge of the surface shows the amount of MU required to offset BTA for a single experiment. From here, the surface slopes sharply down with regard to the number of experiments, showing the great value of experimentation. This is as expected, since the range of services benefiting from end-2-end type architectures grows with more experimentation. In addition, as expected, this growth is at a decreasing rate. The rate of decrease levels out quickly, at around ten experiments, showing that the biggest gain from parallel experimentation is from relatively few experiments.

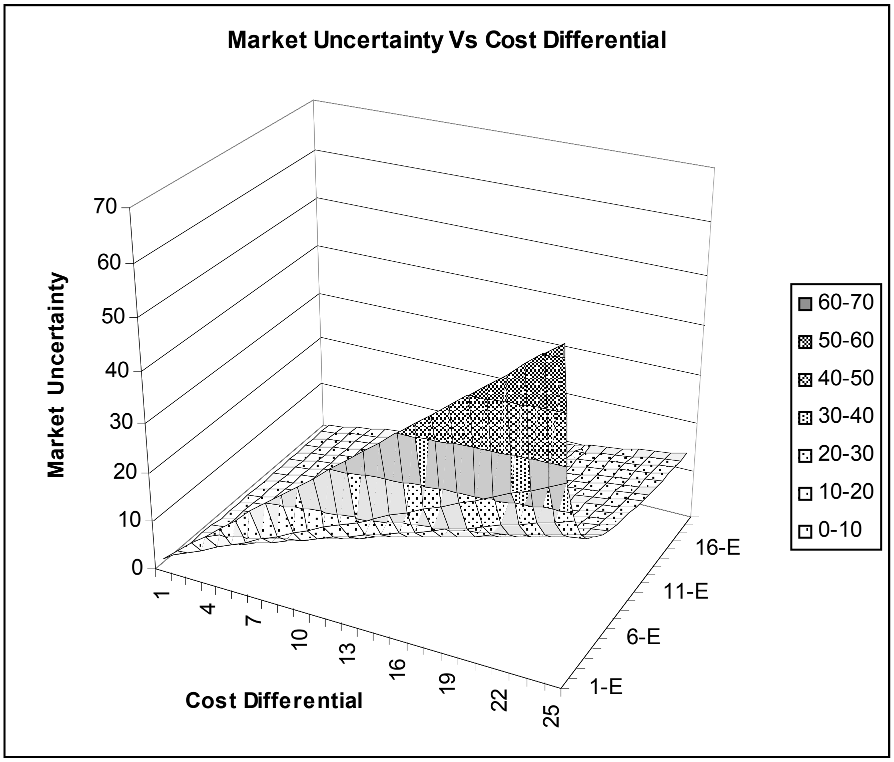

Figure 5 When end-2-end works best

Cost of Modularity

The above arguments show that modularity is good, but it is hard, and can be expensive (as discussed in Section 0) which limits the granularity of modularity in a system. Defining modules that work together and are stable is very difficult. It is expensive to determine if different modules are compatible and will interoperate. Thus, the cost of testing modules and integrating them with the other protocols limits the number and complexity of modules. This cost of modularity is the fixed cost of initially creating the module, plus the cost of testing each changed (or new) module, plus the cost of testing the integration of new modules with the entire system. Baldwin (Baldwin and Clark 1999) shows how to factor this cost of modularity into the benefit of choice.

Staged Development of the Modules

One advantage of modular standards is that they promote staged development. It is easy to add simple modules, and evolve each module independently. This ability of a protocol suite to evolve is important. Dyson says: "*we should not attempt to construct the Internet, but we should act like gardeners, providing a conducive environment for growth*" (Postrel 1998). This suggests that protocol suites that evolve from a simple start will generally achieve a higher value than protocol suites (or individual protocols) that are initially complex. A protocol should be simple, to solve a current but focused problem. Email on the Internet is an example of a service that started out with simple protocols and evolved in complexity. It is far more successful than X.400, the OSI mail protocol that started with many more features[3]. Only after email had established itself, did application protocols for transferring (SMTP) and accessing the email on the local email server (POP and IMAP) become standard Internet protocols. Furthermore, at first, only text-based e-mail was possible; later attachments (via MIME) that allowed binary files as mail become standardized. The first Internet mail specification (RFC 561) is 4 pages long, compared to the current email specification (RFC 2822) which is 51 pages long; the current MIME extensions (RFCs 2046 - 2049) comprise over 100 pages of specifications. This ability to evolve is essential to survival in uncertain environments. Unlike the unsuccessful and over-specified X.400 protocol, Internet email protocols evolved into a set of standards that provided a feature set users wanted, and thus adopted.

One way to value an evolutionary style of enhancement to a protocol stack is to place this evolution in the context of a multi-staged investment. Similar to the example in Amran's *Real Options* (Amram and Kulatilaka 1999) and related work by Kulatilaka (Kulatilaka, Balasubramanian and Storck 1999), the evolution of a protocol stack is viewable as a series of staged investments. Each stage of development creates an option value by providing the choice of whether to continue evolving the stack, and how the protocol suite should change. Staging the investment required to develop a comprehensive set of protocols (or a complex single protocol) minimizes the risk of bad decisions in uncertainty.

Figure 6 shows the first two stages of a hypothetical evolution of protocols in the Internet. Stage one begins with market acceptance of a minimum set of protocols (i.e. TCP/IP). This stage has a single option: invest in a new transport layer protocol (UDP), or not. During stage one, market uncertainty exists - will the market accept the new standard, or not? Stage two begins with four possible branches. At each stage of development, we have a yes/no decision to make; then we roll the dice to see the outcome of our investment choice. This yes or no decision to continue with the protocol suite fits nicely with the binomial model. As the tree unfolds, we see all the various paths the evolution of the protocol can take. The example in (Kulatilaka, Balasubramanian and Storck 1999) discusses how a dynamic programming algorithm back-solves this design tree to determine the optimal choice to make at each investment point. This strategy creates value by increasing the range of outcomes and providing relevant information that can be factored into critical decisions. This provides a higher expected value than the single stage approach, where there are only two possibilities: success of the full-blown protocol stack, or not. With standards, it may be impossible to estimate the changes and choices that will arise in the evolution of a protocol. However, this may be useful as a tool to perform a historical analysis of the particular evolutionary path a protocol has taken.

This evolutionary approach gains from the modularity of each layer. Modularity of the protocol suite allows an additional option benefit at each decision point by permitting choice of the best of many proposed protocols to include in a stack. In effect, protocols developed with an evolutionary style and a layered modularized structure allow a double options benefit. That is, the option of whether to invest in a change, and if so what subsequent change to include from the many available choices.

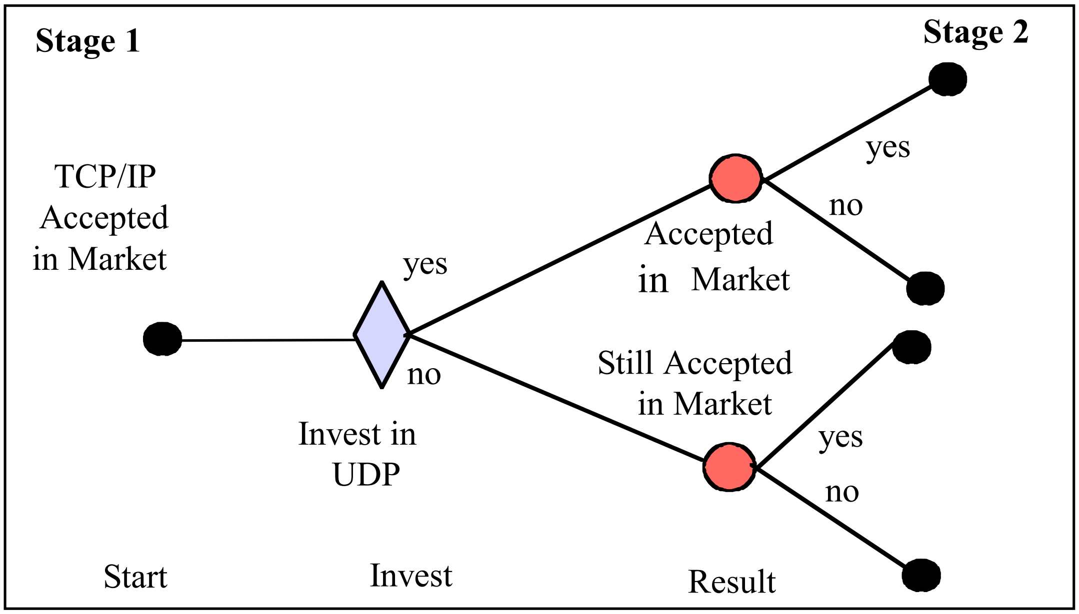

Figure 6 Multi-stage Protocol Introduction

This staged development model is applicable to studying the development of any particular protocol within the stack. The staged approach adds expected value to a protocol's evolution by allowing decisions about whether a protocol's development should continue, and if so, how to alter the protocol at each stage as more current information becomes available. One example of this is how TCP has evolved its sophisticated congestion control scheme. Initially, the congestion control was primitive, but once congestion existed and was better understood, more-effective algorithms became implemented (Jacobsen 1988).

High Level Empirical evidence

Empirical evidence from the Internet supports our model: it evolved as our model would predict. It is modular in design and layered such that it promotes end-2-end services, and the protocols started simple and evolved in complexity along with the users' evolving expectations. The success of a protocol is its adoption by users. For example, as a whole the Internet stack is successful, while the OSI stack is not, but IS-IS is a successful OSI protocol, while Gopher is a failed Internet protocol. Both the IETF and OSI protocols are good examples of modular protocols that allow end-2-end services.

Figure 7 shows the structure of the Internet protocol suite at three points in its history. In (a), we see the TCP/IP Internet stack at its birth, when it consisted of a single network/transport protocol (IP-TCP) and a few applications: ftp (a simple file transfer utility), and telnet (a remote login protocol). We date this at 1974 when Cerf and Kahn (Cerf and Kahn 1974) published their first TCP/IP paper. Then, in (b) we see the next stage of evolution: the separating of IP and TCP into two layers (network and transport) and a new transport protocol (UDP) added for a different type of service to applications. This stage began with a hallway meeting in 1978 (Salus 1995) about the merits of separating TCP and IP. Finally, (c) shows the Internet stack in December 1988 as specified in RFC 1083. This shows the growth of protocols of the Internet stack over a 14-year period.

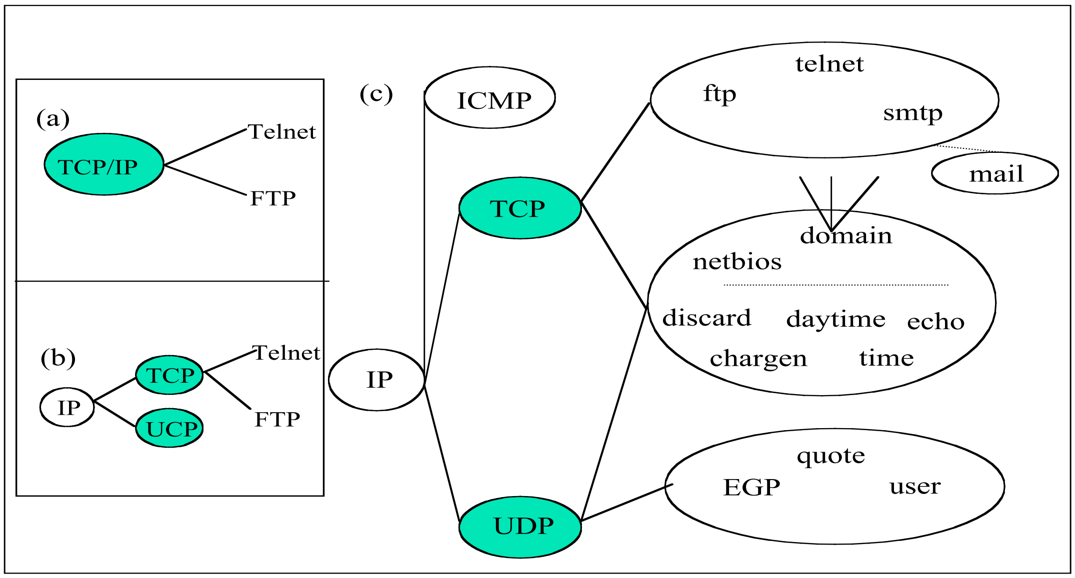

Figure 7 Internet suite structure - History

The standard battle between the Internet protocol stack (promulgated by the IETF) and the OSI stack (standardized by the ISO) is a good example to show the success of technologies introduced as simple standards with few protocols that were then allowed to grow in complexity. The Internet suite of protocols started with only the basic building blocks for a network infrastructure, and these blocks had few options. In contrast, the OSI stack included protocols and options to the protocols, to satisfy every possible need the designers could imagine. Initially introducing standards as "here is everything you ever need" (as in the introduction of the OSI suite) requires over-standardization. Given the high probability of wrong predictions, the kitchen sink approach to standardization is very expensive. One example of this is the five-transport level protocols in the OSI stack, compared to the two in the Internet suite. The Internet protocol suite has shown the diversity of applications possible with the two extremes of the OSI suite: TP-0 is similar to UDP and TP-5, a TCP-like protocol. The OSI argument that different transport protocols are needed to efficiently handle networks of differing reliability turned out to be untrue (Clark, Jacobsen, Romkey and Salwen 1989). Clark shows that the overhead a properly implemented heavyweight transport protocol such as TCP incurs when packets arrive in order without data errors is roughly 234 machine instructions. The market has spoken: OSI is dead for transport, and the Internet stack is the winner. However, this surprised many experts including Marshall Rose (Rose 1990) who believed differently as late as the early 90’s.

Another comparison to make is between the ITU's Frame-relay (introduced as a simple protocol with a successful evolution) WAN protocol and ATM. Frame-relay is an easy-to-understand WAN fast packet switching protocol that began as a short and simple specification from the ITU [I.122 (ANSI T1S1/88-224R)]. Frame-relay originally intended as a layer 2 Common Bearer service for ISDN, is seeing far more success in the WAN marketplace where it has met a well-defined need - connecting LANs over a wide area. Frame-relay is versatile; it can be used to implement a private corporate network, or as an interface standard to connect to a public Frame-relay network. One competitor to Frame-relay is ATM - the Holy Grail of networking protocols. ATM has it all: ultra-fast bandwidth, fine-grained QoS, a complete solution from the desktop to the core. Unfortunately for ATM vendors, users did not adopt ATM as a seamless solution starting at the desktop.

Even with Frame-relay's success, current use is far different than intended by the original providers of the service. At first Frame-relay was unavailable as a lower speed (64K) service, however, now more than 50% of the Frame-relay connections are 64K or less, as a 1999 survey by the Frame-relay Forum shows. Furthermore, never-imagined applications like voice over Frame-relay are evolving. Again, this shows that vendors cannot predict the demands of the users, or their willingness to pay for a service.

Real Options Applied to Protocol and Other IT Standards

Providing standardization options for different contingencies can help reduce the tremendous risk and associated costs of bad decisions regarding standardization. Below we present a generic methodology that, if applied to standards, tends to create a standardization environment that allows the broadest range of experimenters to propose new standards and build services based on them. Our methods promote end-2-end services, which allow more experimentation because anybody including users can innovate. Next, the market selects among the proposed standards, promulgating the standards most likely to be successful. Our theory is simple, intuitive in nature, quantifiable, and empirically verifiable given the success of the Internet.

Our model depends on selection of standards from a market, but this market may be a set of vendors negotiating a standard. As discussed in Section 0, market selection, negotiated agreement, and hybrids between the two are selection mechanisms for standards. Even negotiated standards have an element of market selection because the market can accept or reject what vendors agree to. OSI is a good example of this; vendors and governments negotiated the OSI specification within the guidance of the ISO, but users did not select it - instead they chose the Internet suite. In negotiated standards, one can view the organization responsible for the negotiations to be a market, selecting the fittest technology in regards to vendors, but the user market has the right to reject the agreed-upon standard. Our theory works with both types of selection; it only requires multiple standards picked by users, or agreed to by vendors.

General model assumptions

Since not all standards need to evolve quickly, or be modular in structure with an end-2-end structure, we present a set of conditions that, when met by a technology, implies that standardization of the technology will benefit from our methodology.

1. The conditions within which standardization occurs must have market uncertainty as described in Section 3[4].
2. Technology used to implement the standard changes in predictable ways, but rapidly with short life cycles. The short life cycle of the technology requires flexibility and timeliness in creating standards.
3. The market for services enabled by the technology exists, meaning that profit-seeking firms will have incentive to provide the service, since customers are willing to pay for it, even in the absence of a standard. Without a market, market selection is not possible; the standard becomes anticipatory. While the general market exists, the particulars are uncertain. The particular feature set that will fit best with the market is unknown (i.e. email).

Standardization Rules

Below are three general rules that help limit the risk of standardization of technology with market or other types of uncertainty, based on our real options approach to standardization. The technology standardized should meet the general assumptions given in Section 0.

1. Standards should have a modularized architecture structured to allow the broadest range of experimentation in terms of number of experiments and groups able to contribute. That is, they must allow users to experiment with end-2-end applications. The standardization process must allow market selection of the best outcomes.
2. A good way to introduce standards is in an evolutionary fashion - starting out simple and building the complexity as MU decreases - thus allowing staged investment in creating and growing the standard.
3. Implementing a proposed standard is a good way to show it is possible. Furthermore, multiple independent implementations of the standard help show its clarity, completeness, and interoperability.

Conclusion

We have put forth a paradigm for standardization based on modularity and evolution of complexity when there is uncertainty in the market. Supporting our view is a real options-like model and empirical evidence. Our work shows the value of a layered modularized protocol architecture that allows end users to build services, is simple at its initial introduction, and then evolves in stages as market demands chart their chaotic path. The modular design of a protocol suite provides additional value by giving the designer a portfolio of options of many protocols from which to choose. This is more valuable than a single option on a single complex protocol with many functions. Introducing a protocol that solves a focused problem, and then extending this protocol in stages, maximizes the expected value. Furthermore, we show the value of the end-2-end argument because it allows users to innovate. Our quantitative prescriptive model is intuitively obvious and fits the empirical evidence of the Internet protocol suite and its early standardization process.

Acknowledgments

Thanks to Carliss Baldwin, Barbara Feinberg, and Alan MacCormack - all from Harvard Business School - for many great ideas and brainstorming sessions; Carl Cargill from Sun MicroSystems for encouragement; H.T. Kung (Ph.D. advisor) and Marco Iansiti (of HBS) for supporting this research; great ideas, and motivation, and editorial support from Mark's wife Gretchen Gaynor, and further editorial assistance from Liz Pennell.

---

[1]This market may be users, or a group of interested parties (such as the IETF), as discussed in Section 0.

[2] This is not true from the viewpoint of the service provider. For them it is more valuable to use an architecture where experimentation is possible by the service provider.

[3] This is particularly interesting given that X.400 will run on the Internet (via RFC 1006), but still did not become popular.

[4] This market uncertainty may be unknown by the vendors, for example, with ATM the players were very sure they were right, but were very wrong because of market uncertainty.

References

[1] Gaynor, M. and Bradner, S. (2001). The Real Options Approach to Standardization. *Proceedings of Hawaii International Conference on Systems Science*.

[2] Clark, K. (1985). The interaction of design hierarchies and market concepts in technological evolution. *Research Policy,* 14, 235-251.

[3] Utterback, J. and Suarez, F. (1993). Innovation, competition and industry structure, *Research Policy,* 22.

[4] Dosi, G. Technological paradigms and technological trajectories. (1982) *Research Policy,* 11, 147-162.

[5] Nelson, R. and Winter, S. (1977). In search of useful theory of innovation. *Research Policy* 6, 36-76.

[6] Anderson, P. and Tushman, M. (1990). Technological Discontinuities and Dominant Designs: A cyclical model of technological change. *Administrative Science Quarterly*, 35.

[7] Abernathy, W. and Clark, K. (1985). Innovation: Mapping the winds of creative destruction. *Research Policy* 14.

[8] Brown, S. and Eisenhardt, K. (1997). The Rate of Continuous Change: Linking Complexity Theory and Time-paced Evolution in Relentlessly Shifting Organizations. *Administrative Science Quarterly*, 42, 1-37.

[9] Gould, S. (1992). *Bully for Brontosaurus*. New York, NY: Norton & Company.

[10]Iansiti, M. (Fall 1995). Shooting the Rapids: managing product development in turbulent environments. *California Management Review*, 38 (1).

[11]Rhoden, D. (June 1999). Real-Time On-Line Standards - The Ever Changing JEDEC and Other Standard Topics. *Technology Standards & Standardization Processes, Their Impact on Innovation & International Competitiveness*: *Proceedings of a Lecture Series Presented by US-Japan Technology Management Center*, Stanford University.

[12]Cargill, C. (1997). *Open Systems Standardization, A Business Approach*. New York, NY: Prentice Hall.

[13]Vercoulen, F. and Van Wegberg, M. (2000). Standard Section Modes in Dynamic, Complex Industries: Creating Hybrids between Market Selection and Negotiated Selection of Standards. Working Paper*,* Eindhoven University of Technology, School of Technology Management, The Netherlands.

*[14]*Amram, M. and Kulatilaka, N. (1999). *Real Options, Managing Strategic Investment in an Uncertain World*. Boston, MA: Harvard Business School Press.

[15]P. Balasubramanian, Nalin Kulatilaka and John Storck, Managing Information Technology Investments Using a Real-Options Approach*, Journal of Strategic Information Systems, Volume 9, Issue 1 (1999)*

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

[17]Merton, R. (1992) *Continuous-Time Finance*. Oxford, UK: Blackwell Publishing.

[18]Taudes, A., Feurstein, M., and Mild, A. (2000). Options Analysis of Software Platform Decisions: A Case Study. *MIS Quarterly*, 24(2).

[19]Taudes, A., Feurstein, M., and Mild, A. (1999). How Option Thinking can Improve Software Platform Decisions. *Working Paper Series: Adaptive Information Systems and Management in Economics and Management*, Working Paper 38, www.wu-wien.ac.at/am/Download/wp38.ps

[20]Holland, J. (1992). *Adaptation in Natural and Artificial Systems*. Cambridge, MA: MIT Press.

[21]Kauffman, S. (1993) *At Home In The Universe: The Search for Laws of Self Organization and Complexity*. Oxford, UK: University Press.

[22]Waldrop, M. (1992). *Complexity*. New York, NY: Simon & Schuster.

[23] Hiller, F. S., and Lieberman, G. J. (1967). Operations Research. Holden-Day Inc, San Francisco, p. 631.

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

[25] MacCormack, A. (2000). Towards a Contingent Model of the New Product Development Process: A Comparative Empirical Study. *Working Paper 00-77*, 2000, Harvard Business School, Division of Research.

[26]MacCormack, A. and Verganti, R. (2001). Managing the Sources of Uncertainty: Matching Process and Context in New Product Development. *Working Paper 00-78*, Harvard Business School, Division of Research.

[27]Saltzer, J., Reed, D., and Clark, D. (1984). End-To-End Arguments in system design. *ACM Transactions in Computer Systems* 2(4), 277-288.

[28]Berners-Lee, T. (1999). *Weaving the Web*. San Francisco, CA: Harper Publishing.

[29]Cerf, V. and Kahn, R. (1974). A Protocol for Packet Network Interconnection. *IEEE Transactions of Communications*, (vol.).

[30]Salus, P. (1995). *Casting the Net, From ARPANET to INTERNET and beyond…* Boston, MA: Addison-Wesley.

[31]Hippel, E. (1998). Economics of Product Development by Users: The Impact of "Sticky" Local Information. *Management Science*, 44(5).

[32]Gaynor, M. (2001). The effect of market uncertainty on the management structure of network based services. Ph.D Thesis, Harvard University.

[33]Postrel, V. (1998). *The Future and its Enemies.* New York, NY: The Free Press.

[34]Jacobson, V. (1988). Congestion Avoidance and Control. *Proceedings of the SIGCOMM'88 Workshop*, 314-329. ACM SIGCOMM, ACM Press.

[35]Clark, D., Jacobson, V., Romkey, J., and Salwen, H. (1989). An analysis of TCP processing overhead. *IEEE Communications Magazine*. (vol).

[36]Rose, M. (1990). *The Open Book: A Practical Perspective on OSI*. Saddle River, NJ: Prentice Hall.