Internet Architectural and Policy Implications[1]

for

Migration from High End User to Low End User

Terrence P. McGarty
The Telmarc Group, Inc.

Carole Haywood
RAM Mobile Data

Abstract

The Internet has been focused on a class of users who have been centered around Universities and Research establishments. Recently the user community has been broadening into the less sophisticated technical user, one who is interested in what Internet provides and less in Internet as a tool in their research. This paper presents a focus on that new and expending community and discusses its structure, its needs, the means of accessing that user group and the policy issues relating to this new segment. This paper considers the technical, economic, and policy implications of extending Internet access to the low end user in the broadest context. This will imply ISDN and higher data rate interfaces accessible to the end user community, using such modes as CATV and wireless. It raises the questions of end user access fee costs, the ability to have a Universal Service Option, the modes of access and the architectural evolution of Internet into this community.

1.0 Introduction

Recently a Cable Television company, Continental Cable, and PSI, a commercial Internet carrier announced access to the Internet on a full time basis for about $100 per month to any of the CATV company's customers in Cambridge, MA. The response was overwhelming, and the customers are now awaiting the CATV company to activate the service offering that they promoted. In the same time frame, AT&T has started a consumer trial in New Jersey to give residential users access to Internet for a free trial period to determine what uses they may put it to. One of the users is a group of antique appraisers who have found that Internet allows them to share information, compare appraisals, and communicate with professionals on the provenance of certain rare items. These events define the growth and development of the low end Internet community. The low end Internet user is the consumer user, the casual user, the local library user, or the secondary or primary school user. The low end user is a user of Internet initially and ancillary to their general needs, but it frequently becomes an integral part of what they do. The low end Internet user is that latent user who may in the long run become the mainstay of the Internet community.

In so doing, we first focus on the existing architecture of the Internet and from that discuss its evolution in three dimensions. These dimensions are perceived as the main drivers for expansion into the end user community. The technology drivers are multimedia communications, increased access alternatives, and host migration from the location to the individual. It is essential to determine and define the underlying assumptions and paradigms with which the Internet operates. This paper defines these in some detail, using the historical evolution as a prelude to the future development. Specifically we as; ***what are the current architectural constraints of Internet and how may these be modified to ensure a full and complete development for the low end user community?***

We then focus on the needs and opportunities for this group of low end users. We develop an analysis of the low end user community and show how it differs from the current Internet community. We then discuss

[1] Presented at the Kennedy School, **Harvard University**, May 26, 1993, at *Public Access to the Internet*, Cambridge, MA.

the uses that the low end users may put the Internet to in their applications. The user community is a growing and expanding base, and any discussion we make here may most likely be modified dramatically as new users enter the community. The trend, however, is projectable; namely, that the Internet community will expand in new and innovative ways to users now unlikely to be users. The key question we address is; ***who will be the low end users and what demands will they place on Internet, from both a technical and applications perspective?***

The evolution of the Internet into the broad low end user community will require changes to the current architecture in three areas; multimedia communications, access expansion and host migration. These three technological and architectural changes are discussed in detail. Multimedia communications is in its infancy. We present several of the key challenges that Internet faces if it were to provide these capabilities to the low end user. Access is the second, but in the short term, the most driving issue. Access is an issue of cost combined with flexibility. It is the issue of through what means does the low end user obtain real time communications access to the Internet. The last issue is host migration. It begs the question of whether a Personal Digital Assistant is or may become a host on the Internet. This then leads us to focus on the issue; ***what are the driving technologies to increase access to and applications of the resources available to the low end user community, and what can the current Internet community do to facilitate these?***

Having developed a set of constructs for the Internet in terms of its evolution, we then address the issues of the economic structure of the Internet as it moves to address the needs of the low end user. These needs are driven by access and pricing, and pricing must have a rational basis in terms of the overall Internet structure. As such, we then develop a construct for modeling the cost structure of the Internet and then apply this model to an overall microeconomics model of the Internet. We show that there is, as of this time, inadequate information on the costs structure of the Internet to provide a solid base for pricing. This leads us to question; ***how do we develop a proper costing and pricing model for Internet so as to allow for a balance of price for access with revenue to cover costs?***

We conclude the paper with a discussion of the policy issues that are brought out in developing the overall evolutionary alternatives for the low end user migration. The authors believe that the migration is doable and can be readily achieved in the context of the current design and design methodology. The policy challenge is to answer the questions of Internet evolution; ***namely, is Internet a scaleable economic entity that creates value that will be accessible to the widest possible collection of users, while retaining the value already provided to the high end user community?***

2.0 Current Architectural Elements

The current architectural structure of the Internet has been described elsewhere at length. We review it briefly here to establish a basis for a detailed description of the system.

2.1 Internet Evolution

Before doing so it is important to view the evolution of the INTERNET concept in the context of several Phases. These phases are:

Phase 1, The Simple Internet (1968-1974): Beginning as an experiment in networking, 'the ultimate petri dish', in this period the Internet, as ARPANET was a simple interconnection of at most 56 Kbps circuits interconnected by Intermediate Message Processors, IMPs. The user community was a collection of large scale computer processing facilities with end users identified to their hosts. The concept of time sharing in an operational context was not present. The original thoughts of INTERNET users use was access to complex computer resources from afar. The typical example was access to an ILIAC IV from a user at Berkeley.

Phase 2, INTERNET Goes Global (1973-1981): In this period the TCP/IP protocol is developed and added on top of the existing datagram network. Although originally aimed at remote login and file transfer (FTP), the afterthought, Email, become 95% or more of the total network traffic. Email users now expand their identity from host to self attached to host. The INTERNET goes international with satellite connections to Goonhilly in the U.K. COMSAT is used as a access node to the European community. Initially the international use is justified by the seismic data transfer needs and the requirement to backhaul large volumes of traffic. Ultimately Email still dominates in all areas.

Phase 3, Military and Non Military Split (1982-1986): During this phase the non-military INTERNET evolves. DoD separates its network and the residual is spun off into a larger user community. Unauthorized accesses are found and user scares from hacking is observed. The user community expands allowing access to multiple user communities. The backbone grows to T1 rate, namely 1.544 Mbps.

Phase 4, The Mitotic Period (1986-1992): "Cell" division of the network occurs. DS3 or 45 Mbps circuits are added and local and regional networks are adopted. The proliferation of access closer to the end user is generated and hosts grow explosively. Networks seen by the INTERNET grow from about 100 in 1988 to over 5,000 in 1992. Personal computers proliferate, and access to the end user is now growing. Identity is still with the host. The dominant traffic is still Email and it derivatives, and the first virus, a self replicating 'worm' ran amok for several days in 1988.

Phase 5, Low End User Access Era (1993-Present): This era is the era of low end user access and the proliferation of commercial user hosts and networks. The user community is expanding from the computer literate and comfortable to the infrequent user community.

Phase 6, The Distributed Open Network (1996-) : The network moves into a Giga bit per second backbone allowing for the first time real time access. The protocols for access allow expansive addressing and accessibility. End user access costs are reduced by access costs enablement/control policies and the introduction of 64 Kbps end user transport access to all terminals. Host identity is now made consistent with user identity.

2.2 Architecture

The current Internet is structured as shown in Figure 2.1 as a four layered structure. There are the backbone, the Regional, and the Campus networks and the Hosts.

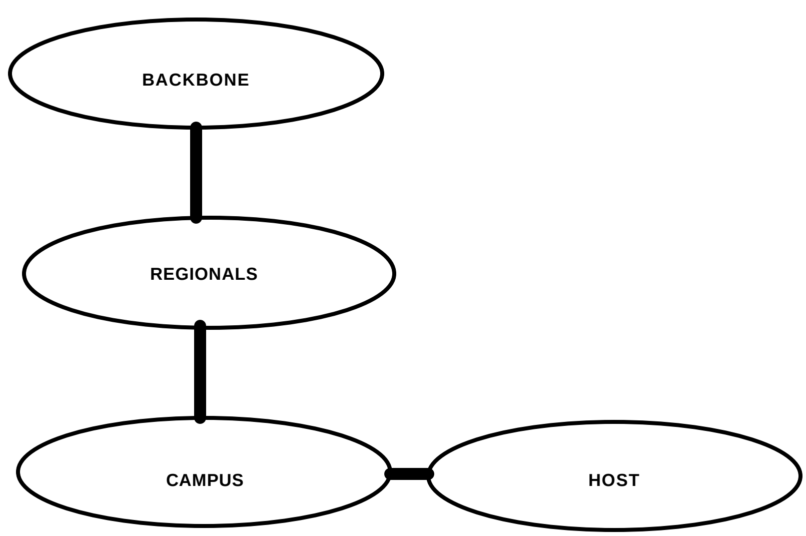

The essence of Internet facilitation and accessibility is the set of protocols available to the community to allow access by a wide variety of hosts in a complex and fully distributed fashion. The protocols are at the heart of Internet success. They are the "software and system agreements" that allow disparate machines and software to talk across equally disparate networks. The current protocols focus on data transactions, with some innovation allowing images and limited multimedia; namely voice and video. The future challenge will be the development of new and innovative protocols to allow both low end user access to grow while at the same time enriching the capability of the information transferred.

The key underlying protocol structure that makes the Internet function is the Transport Control Protocol/Internet Protocol, TCP/IP protocol suite. This protocol allows for the easy and ready flow of data from one user to another by agreements at various levels of the network to handle, process, manage, and control the underlying data packets. Protocols such as TCP/IP will be the heart of the evolution of the Internet. We shall focus latter on such protocols as applied to multimedia and new access methods. One can best understand the protocol evolution by looking more closely at TCP/IP. To quote Cerf:

> *"IP (the Internet Protocol) provides for the carriage of datagrams from a source hosts to destination hosts, possibly passing through one or more routers and networks in the process. A datagram is a finite length packet of bits containing a header and a payload. ... Both hosts and routers in an Internet are involved in processing the IP headers. The hosts must create them ... and the routers must examine them for the purpose of making routing decisions, and modify them as the IP packets make their way from the source to the destination."*[2]
>
> *"TCP is a protocol designed ... to provide its clients at a higher layers of protocol a reliable, sequenced, flow controlled end to end octet stream...."*

The development of new protocols can best be determined from studying the twenty year evolution of the TCP/IP protocol. The rationale for many of the TCP mechanisms can be understood through the following observations:

> ***1. TCP operates above IP and IP provides only best efforts datagram transmission service.***
>
> ***2. End to end recovery... leads to sequencing..***
>
> ***3. Flow control requires that both ends uniquely agree...***
>
> ***4. In a concatenation... it is possible for a packet to circulate...***
>
> ***5. Termination ... should be graceful..***
>
> ***6. Every process should be able to engage in multiple conversations***
>
> ***7. ... the arrival of information should contain no semantic differences..."***[3]

---

[2]See Cerf, pp. 84-85 in Lynch and Rose.

[3]Cerf, pp. 117-118, in Lynch and Rose.

These types of architectural and system requirements must be articulated as clearly and carefully for each of the new dimensions of expansion of the Internet. TCP/IP protocols have emerged as the standard network interface to the Internet that allows users to send messages from one point to another in a reliable form and also the users to embody in those messages certain characteristics that make them more then just a collection of bits. IP gets the packet across the network, and TCP brings the underlying nature of the packet stream into context as a reconstituted entity.

The current Internet Architecture thus has two main elements. The first is the semi hierarchical structure of Backbone, Regional, Campus and Host, and the second is the agreement on a single protocol to talk across the Internet in the TCP/IP suite. These elements reflect a great deal about how the Internet is managed, and what its growth potential truly is. Several observations can be made about the Internet in its current embodiment:

- **Host Orientation:** The Internet is host oriented. It is focused around the host as the terminal entity. This will imply that the concept of a host may have to be expanded to include a new and wider variety of electronic entities, some physical and other actually virtual devices. Concerns about wireless access devices, Personal Digital Assistants, distributed processing devices must be considered.

- **High disaggregation of data:** The network assumes a high degree of disaggregation of the data from one location to another. The current assumptions are that data from one location are independent of data from other locations. In a multimedia environment, this will no longer be the case. Data will be virtually aggregated into a compound multimedia object, thus creating a virtual aggregation on a disparate spatial disjoint set of data elements. Specifically, my mouse movement at one location will be related to my voice at another, and a third parties video at a third. The concatenation and orchestration of these disparate entities will be viewed as a single totality.

- **Poor Transmission Link Performance:** The underlying structure of IP was built to deal with a poor quality transmission path, not those found in fiber optic networks. The new fiber optic networks are almost error free In fact, with the processing at higher protocol layers, one may now assume error free transmission. This means that latency prevalent from the old analog telephone lines may be totally eliminated. This is a critical factor for multimedia transactions. On the other hand, the new wireless communications services may have higher data rates until one better understands how to deal with multipath and other radio propagation factors. Thus the network must handle error free fiber, combined with error prone wireless.

- **Low Intra Network Intelligence:** Limitations of processing in the network are due to the simplicity of the routers and the "intelligence" of the host. New network elements are highly intelligent, and even the PDAs and the wireless devices contain dramatically greater intelligence to perform processing. The ability of the network to drive more processing to the periphery and to the new fully distributed host environment will enable the Internet to added new degrees of both access and services flexibility.

- **Moderate Speed Transport:** The layered structure of the protocol suite allows for flexibility in low to moderate speed networks. At higher speeds the protocol suites begin to breakdown. Fiber networks will rapidly allow Gigabit per second transmission, and high resolution images and video will be integrated in complex multimedia objects. This will require a significant rethinking of what the object control must look like and the development of new and innovative protocols.

- **Single Media Messaging:** The TCP layer focuses on getting a single stream of data through. Modifications can be made for voice or even video streaming but a full multimedia network is not achievable. An enhanced multimedia TCP/IP type system will be constructed that allows the entire suite of users access to multimedia sessioning with high data rates and accessible via fully distributed high end processing devices, albeit at dramatically lower cost.

Internet was intended as at data communications tool across the academic community. No intelligence was built in to the network to guarantee arrival of messages, nor to mitigate congestion. Issues of security and virus protection are management options left to the discretion of the host configuration.

We see from the above set of observations and trends that these architectural assumptions are being challenged, from both the high end users as well as from the low end user. The Internet backbone is upgraded to 45 Mbps, DS 3, circuits, and the campus networks are upgraded to 100 Mbps and above networks. The Regionals are mainly at 1.5 Mbps, DS 1, rates. However, fiber migration, CATV migration, and even wireless will change this. Multimedia communications and wide spread access for low end users will also challenge the needs for more Intelligence in the network, not just layered on the protocol.

3.0 The Low End User Community

The Low End User Community is potentially the fastest growing group of Intemet users. In this section we discuss the low end user community and describe some of the uses that they can make of the Internet. Then we discuss several of the key factors that must be part of these considerations. Specifically, what is the information does the user have access to and how do we value that information. We argue that the concept of value is critical to understanding whether the use of the Internet is more than a passing fancy or a fundamental structural change both in society, and the way that members of society communicate with one another. All too often, we see approaches that list hundreds of uses with no understanding or appreciation for what must be altered or changed in the overall infrastructure to accomplish the change and more importantly with no understanding of the structural value of the change.

3.1 Users and Uses

This section presents a set of five user communities which can be considered low end, and briefly discusses their uses. These are but a few of the many that may evolve but they represent the most visible in terms of current acceptance and value.

(i) **Commercial Casual Users:** The casual user is the user who has recognized that Internet provides a community of interest and access to information in its broadest base. The commercial casual user is the one in a corporate environment whose initial desire is to use Internet for Email and other such limited applications. If one recalls the ARPA Net in the mid 1970s, the dominant use was Email, being well in excess of 95% of all use for ARPA Net users. Now Email is below 30% and dropping as other uses are discovered. The Email functionality is an entry point into the Internet. It is the service that will capture the casual commercial user, say an R&D Manager or Executive, whose staff of researchers are already high volume users of Internet.

(ii) **K-12 user Community:** The K-12 is an explosive area. There are examples in many districts of schools at all levels encouraging students to use the Internet as both an educational tool as well as a cultural learning device. Internet, as a global network, allows students to communicate with others around the world. The classic story is the student who wrote an 'A' graded paper on Aborigines by communicating with Australia directly over the Internet. There are two countervailing factors, however, that are pressing at this opportunity. One, is the growth of the multimedia curriculum exemplified by in class video education that is real time. The other is the fact that over 98% of classrooms have no telephone access. In most if not all cases of education through new electronic media, it is easier to get the satellite dish, the VCR, the television, and the computer, then to get the RJ-11 telephone jack! We will argue in this paper that CATV and wireless will make a dramatic change in this area. [4]

---

[4] In fact, the Burnaby South Secondary School, which opened in British Columbia in February 1993, has built its visionary 'village of learners' campus around a communications infrastructure of coaxial cable in every room to take advantage of multimedia learning aids. It refers to its imbedded network capabilities as its 'information highway'.

(iii) **Health Care User Community**: The Internet is a data backbone network that is in essence the basis of a multimedia transaction based infrastructure. Internet will eventually allow the transport of images, voice communications and already encourages specially and temporally displaced conversationality. Furthermore, Internet is a transaction network, transacting everything from data transport to Email. Health Care is an industry in search of productivity improvements. It has been shown elsewhere, that over 35% of the expenses in Health Care are due to the handling of paper for patient billing and non critical record management.[5] Also, quality of care, as contrast to cost of care, is improved with increased flow of information and displaced conversationality between not only the specialties but also the total Health Care Team. It is argued that the Internet can be that catalyst that enables the internetting of Health Care for both productivity improvements and quality of care improvements.

(iv) **Higher Education User Community**: Not all Universities have access to the Internet, nor do all students have either the proficiency or knowledge of what the Internet can do for them. In many ways it has been the "techy's" toy that is now expanding its way into other areas. This can bee seen with the number of user groups exploding in many non technical academic areas as well as with the explosive growth of electronic publishing. The later is the first true example of the growth of an electronic industry with the enabling capability of an electronic marketing and distribution channel, namely the Internet.

(v) **Residential Casual Users**: The residential users is becoming one of the largest growing segments. The residential user has had access to such networks as Compuserve over the years and also Prodigy, the erstwhile videotext attempt by Sears and IBM which has never been profitable. Ironically Compuserve, a text driven system, is successful because it meets the consumers needs. The Compuserve user is migrating to the Internet since the Internet is a broader community of users and also because University students can communicate home via the Internet. Thus access to the Internet via Compuserve or direct access via Gateways such as General Videotex, are opening doors for the residential user.

3.2 Pricing

There are two major dimensions that relate to pricing. The first dimension is economic and relates to the asset, its value, and payment for that value (we shall delay discussion as to what the value is). The second dimension is the market dimension of pricing which focuses on the need to stimulate the user demand and to focus on needs and services rather that cost and capital recapture. Let us consider the Market driven pricing issues first and then the Economic driven issues.

Market Driven Pricing: With any new and innovative product or service, users may not perceive the value until they have had the opportunity to experiment with it at length. An attractively and predictably priced service would represent a low barrier of entry to use and encourage users to experiment to experience the benefits, which would soon become a need. Thus pricing can be used as a motivator to users of how best to gain access to the services. We shall argue that this focus of predictability demands a fixed base pricing scheme which is usage independent, based upon class associations. Thus all residential casual users should be charged a fixed fee and unlimited local usage.[6]

Cost Based Driven Pricing: This type of pricing is classic in regulatory circles. It assumes that a model exists for the network, that such a model has a readily assignable allocation procedure to users not just user communities, and that rates of usage and rates of returns can be measured and determined. With these assumptions in place, then a usage sensitive pricing scheme can be implemented. This is in essence what

---

[5]McGarty, Health Care Policy Alternatives, Telmarc Group Report, 93-004, May, 1993.

[6]Solomon has argued this even for Broadband ISDN and for the same set of reasons. In addition, Theodore Vail did this after almost twenty years of measured serviced in the early days of AT&T.

the Congress and the FCC has recently done with Cable Reregulation. The major problem with this form of pricing is that is stifles innovation and more importantly the cost to the consumer of meeting the Government strictures far exceed the gains.[7] There is a second dimension of this type of pricing depending on the availability of capacity. If there is excess capacity, then fixed pricing is possible because no users will be deprived of access. If capacity is limited then variable usage based pricing is the way to go because it auctions off the limited available capacity. Other perturbations of these pricing schemes are available.

Another type of pricing is value based pricing. This takes into account the dynamics of the Market form converging into a set of articulateable needs representative of information value.

Value Driven Pricing: The value based pricing scheme is independent of costs and assumes that the user have already made a value judgment and reflects that in a demand for the service, rather than just reflecting benefits. The demand is predicated upon a structured value judgment that is quantifiable in their use of the service. This means that a possible pricing scheme is based not upon costs but upon value transfer or creation. This scheme is used in financial markets but not in regulated industries. This third pricing option is certainly viable for the unregulated, multi-network value added capabilities accessible through Internet.

We have developed the constructs of value and the constructs of price. This will allow us to develop a demand model for the Internet service. We divide the market into three segments;

Research Community: These are the current set of research users of the Internet. To this user community the Internet has a value based on the users ability to obtain information necessary for their research and to have access to a transaction capability to communicate this with others on the network. It also is a path that enables access to other network attached resources such as supercomputers.

Commercial Community: These are the commercial users who use the Internet for conductivity, communications and computation. They also have other strategic uses that more closely map into the information value alternatives presented earlier.

Low End User Community: This is the nascent community of low end users who are not full academics and researchers nor are they larger commercial users. Low end users are an entity unto themselves. Their value is based upon a more consumer-like demand model.

Now we can create a demand model that assumes a quantity, q, and a price, p. Let us first discuss the quantity and then the price.

Quantity, q: The quantity may be a minute of access, a months worth of unlimited time access, a file access, a storage access, service access, or some other form of access. We shall not describe what that is but it is clear that user response will be fashioned by what the unit of Internet quantity, q, really is. We shall abstract the quantity for the time being, but it will become a key policy issue as we progress through the microeconomic analysis.

Price, p: The price will be based on one or several of the value schemes that we have presented and will follow accordingly one of the three pricing schemes. The key issue will be the relationship of price to what the quantity is as an entity.

The demand curve is;

---

[7] A comparable issue arises in LEC network access. The access fees are barriers to entry to LEC competitors. The establishment of co-carries status is a means of elimination, see McGarty, Wireless Access, MIT, March, 1993. However this results in settlements to have each carrier pay the other the residual difference. However, the cost of measuring and collecting and processing this settlements data far exceeds the settlement amount. Thus it has been argued that full and free Open Access is Pareto Optimal.

p = f(q)

Figure 3.1 depicts the demand curve as three individual demand curves and the total demand curve. This begs the question of how the market will actually measure, create, and motivate demand.

Figure 3.1: Demand Model Construct for Internet

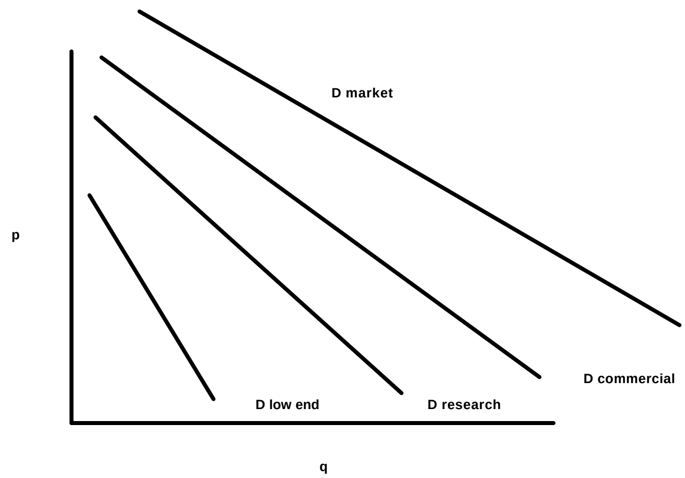

4.0 Architectural Evolution

The low end user community has a different set of requirements for both use of, and access to the Internet facilities. The emerging demand for multimedia services and the potential mobility of the low end user present several challenges to the architecture which may change its elements in an evolutionary sense. In this section we develop some of the key new dimensions in the expansion of the Internet architecture. These dimensions reflect not only the needs of the low end users but also of the remaining established user base. The main catalyst is the low end user however, because the expansion of applications, the reduction in costs of access and the redefinition and commoditization of the host concept are fundamentally driven by the expended nature of the low end user applications.

Local access is one key issue. ISDN was and still is an option, yet it always has been too little and too late. Alternative access mechanisms such as wireless and CATV may lead to the end of ISDN. ATM switching is another element that has a great deal of focus. However, ATM is another hierarchical architectural solution to a problem that may be shifting to a more fully distributed architecture.

A future issue is whether dark fiber providing bandwidth capability coupled with intelligent hosts handling the switching in a highly distributed fashion will become architectural options. Until dark fiber becomes an available commodity, switching will remain a centrally provided host service. Host Intelligence will, in any case, be the driver.

Another key issue centers on the ability to evolve from Client server architecture to a peer to peer protocol set. At higher bandwidths and access rates this may result in a convergence of flexible peer relationships that become applications dependent. In this section we layout several of these options.

Specifically, we believe that there are three main technological changes which will impact low end user migration as well as the existing base. These are;

Multimedia Communications: This area of multimedia communications is generally the least understood and most discussed area in both computers and communications. The challenge of multimedia communications is to create what we have called "displaced conversationality". This means the provision of all sensory inputs and outputs to any human user at any time and place required for the transmission of information in order to transact a series of events, leading ultimately to an agreed consensus amongst the parties involved in the transaction. Simply put, it means the I can talk in simple terms with anybody else, using whatever displays, video, data, voice or other annotations I desire, either simultaneously of at a delayed period of time. This will place significant new demands on the Internet. It begs the question of whether the Internet must now consider raising the level of protocols it supports above just TCP into what we have called the session control protocol suite, SCP. Does Internet evolve into a SCP/TCP/IP network?

Access Expansion: Access implies any and all physical communications means that a user may have to access the Internet, either through the Campus system, the Regional or even the Backbone. Today, we view the access to be achieved via a telephone line or possibly a LAN. In this paper we extend the access in two dimensions; CATV and wireless. CATV access means broadband access even with the systems in place today. This means 50 to 100 Mbps access in wide areas of coverage and this, combined with the advances in multimedia communications complement one another. The second access innovation is wireless access. Specifically the authors describe PCS, Personal Communications Services, the new and innovation access schemes at 1.8 to 2.0 GHz. This access scheme will enable the extensive Host Migration to PDAs, Personal Digital Assistants and the migration of network identity from host to person, and the demands put upon the network to "Find Me!". The cost of access in this new and competitive environment will be of primary importance. We have seen the cost of IEC access decrease by more that 50% since divestiture.

We expect Local Exchange Access to do the same when competition from these areas are introduced. For example, we envision that if access changes are made as discussed, an end user may have access with standards 10 Mbps ETHERNET boards, via a wireless shared Local Area Network, to a server host, at a rate of less than $50 per month for unlimited access time. Compare this to the current rates that exceed $1,000 per month, and one can envision significant low end user expansion. Access expansion simply means the reduction in the cost of local telecommunications access. This is now a key cost element and a concomitant barrier to entry to a wide selection of low end users.

Host Migration: Historically, an Internet user was identified with a Host. The user had access via the host and the user was merely an extension of this host. This made sense when the user requires access to the host for the host shared resources. With the power, increased capabilities and ubiquity of personal computers, migration of identity from the host to the user is more likely. The development of PDAs or Personal Digital Assistants, which are now user "resident" hosts, rather than host "resident" users, are a driven technological change which will cause significant architectural change in the Internet. The user can now be in possession of the Host and the host can be connected to the Internet in a wireless fashion, thus the need for "Find Me!" functions in the Internet fabric.

We now develop each of these areas and provided some of the fabric for the challenges to the Internet architectural evolution.

4.1 Multimedia Communications

To achieve a truly visionary statement of policy and architectural implications, the definition that we take of multimedia in this paper is an expansive one, not the one that most current technical and marketing experts view as multimedia. The reader is strongly advised to place prior biases about their own understanding of multimedia aside and understand what we define as multimedia communications, a more restrictive definition but a more expansive application.

It has been argued elsewhere (see McGarty [1]) that multimedia should not be confined to merely the storage of information of multiple storage devices. Tomorrow's multimedia communications will encompass processing targeted at all the senses, beyond today's focus on sight, sound and touch, and will include all of the interfaces between humans, processors, storage devices and the network. In fact, for the purpose of this paper we define multimedia as the confluence of storage, senses and interfaces, taking the multimedia paradigm and adding multiple human elements, thereby transcending the prototypical computer communications view of the world.

Multimedia communications is simply displaced conversationality. From the low end user perspective it represents a truly transparent medium of talking, of sharing ideas and conversations with other in a simple fashion, blending seamlessly all elements of communications that typically are in such a true conversation. Multimedia communications is not merely the devices and displays, it is not merely advanced CD players with enhanced sound. It is a conversation with others using all of the available senses, combining meaning and content between a group of individuals, displaced in time and space. It is his vision of multimedia communications that the authors wish to bring to the field of the Internet and especially to leverage the growth and utility to the low end users.

When we introduce the communications concepts, we do so in the context of having multiple users share in the use of the multimedia objects. Thus multimedia communications requires that multiple human users have sensory interfaces to multiple versions of complex objects stored on multiple storage media. In contrast to data communications in the computer domain, where humans are a secondary after thought, and optimization is made in accordance with the machine to machine connection, multimedia a communications is a human to many other human communications process that must fully integrate the end user into the environment. Multimedia communications thus generates a sense of conversationality, it is sustainable over longer periods, and has an extreme fluidity of interaction.

Today's broadband architectures provide higher speed communications but retain the stricture of the 64 Kbps voice paradigm and its hierarchical switching in the world of data packets. Techniques such as ATM and SMDS, as well as FDDI are direct offshoots of this paradigm. These architectures fail to accommodate the paradigm that we are developing in this paper that relates to the structure of the multimedia object and the conversationality of multimedia communications. Existing telecommunications architectures significantly reduce performance and quality, while having high transport costs.

In this paper, we concentrate on three issues in the area of multimedia communications; the data objects, the conversationality of the interaction and the overall communications architecture. It is this concept of conversationality, of displaced sharing of information and creating transactions of the mind that make this definition of multimedia communications more expansive. It is this concept of displaced conversationality that demands interaction and such ***transparent interaction*** is the inherent capability and functionality of the Internet, especially as it addresses the low end user

We first note that the data structures in multimedia environments are dramatically different than those in normal computer data communications. Specifically, Mullender has shown that typical data file sizes that are transferred in a UNIX environment are on the order of 2K bits whereas in a multimedia environment the file size may average 100 Mbits.

Secondly, a multimedia environment needs to handle real time, no latency interactions such as that in real time voice and video. As is well known, such transport protocols as TCP/IP are not adequate from a latency perspective to support these types of information objects. This conversationality aspect of the multimedia environment is key to effective communications. In this paper we focus on utilizing the Session layer from the OSI format for the delivery of the multi-user conversationality. Historically, session layer (See Tannenbaum) has been relegated to a secondary position in the OSI hierarchy. In a multimedia environment, we show that the session functionality, refined and expanded, provides the essential integrating capability for conversationality.

The remaining communications services, at OSI layer 4 and below, become at best, delimiting factors in the communications environment. There are certain underlying performance factors of the lower four layers, that when combined control the overall end to end performance as viewed from the users perspective.

The major observation that can be made is that the standard approach to communications system design, from the physical layer and up is the wrong way to proceed for multimedia. Specifically, in a multimedia environment, one must, perforce of user acceptance, design the system from the top layers and down. We briefly review two of the key elements of multimedia communications that relate to Internet expansion; multimedia data objects and multimedia communications protocols.

(i) Multimedia Data Objects:

In a more standard computer communications environment, the data objects have significant structure and they are frequently integrated into a system wide data base management system that ensures the overall integrity of the data structures. In a multimedia environment the data elements are more complex, taking the form of video, voice, text, images and may be real time in nature or can be gathered from a stored environment. More importantly, the separate data objects may combined into more complex forms so that the users may want to create new objects by concatenating several simpler objects into a complex whole. Thus we can conceive of a set of three objects composed of an image, a voice annotation and a pointer motion such as a mouse movement. The combination of all three of these can also be viewed as a single identifiable multimedia object.

Before commencing on the issues of communications, it is necessary to understand the data objects that are to be communicated. The objects can be simple, or combined into more complex objects. We call the initial objects Simple Multimedia Objects (SMOs) and the combination of several a Compound Multimedia Object (CMO). In general a multimedia communications process involves one or several SMOs and possibly several CMOs.

Figure 4.1 SMO Structure

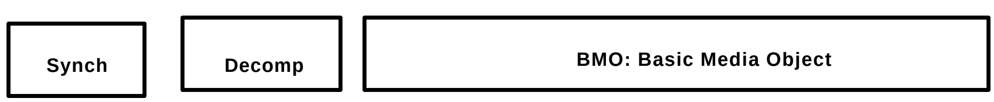

A simple multimedia object, SMO, as depicted in Figure 4.1 is a data string we call a Basic Multimedia Object (BMO), with two additional header fields; a Synchronization field and a Decomposition field. The Synch field details the inherent internal timing information relative to the BMO. For example it may contain the information on the sample rate, the sample density and the other internal temporal structure of the object. It will be a useful field in the overall end to end timing in the network.

There may be two types of BMOs. The first type we call a segmented BMO or SG:BMO. It has a definite length in data bits and may result from either a stored data record or from a generated record that has a natural data length such as a single image screen or text record.

The second type of BMO is a streamed BMO, ST:BMO. This BMO has an a priori undetermined duration. Thus it may be a real time voice or video segment.

When we combine these objects together we can create a compound multimedia object. This is shown in Figure 4.2. A CMO has two headers, the Orchestration header and the Concatenation header. The Orchestration header describes the temporal relationship between the SMOs and ensures that they are not only individually synchronized but also they are jointly orchestrated. The orchestration concept has also been introduced by Nicolaou. In this paper we further extend the orchestration function beyond that of

Nicolaou. The concatenation function provides a description of the logical and spatial relationships amongst the SMOs.

Figure 4.2 CMO Structure

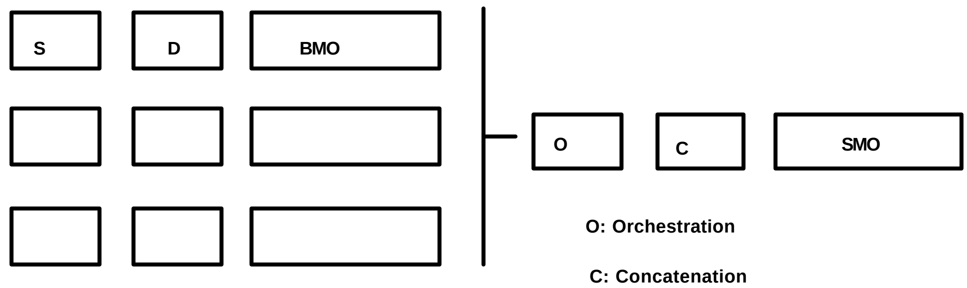

These concepts have been further developed in McGarty[2] and there we have provided more detailed structure to the multimedia data objects.

We can also expand the concept of a CMO as one that is created and managed by multiple users at multiple locations. Figure 4.3 depicts the structure of a distributed and shared CMO. In this construct we have demonstrated that n users can create a CMO by entering multiple SMOs into the overall CMO structure. Simply put, multiple low end users can create a conversation composed of video, voice, text, pointers, all interacting in a single session. Thus a class of students can interact with another class of students and describe a set of physical phenomenon, create a new collaborative piece of art work, or even create a new construct in interactive music, called the distributed symphony. Figure 4.3 depicts the concept of the session having a concatenation of multimedia objects, that the network, at higher protocol layers must transparently manage these and that new and innovative protocols are at the heart of this innovation.

Figure 4.3 Distributed CMO Synchronization

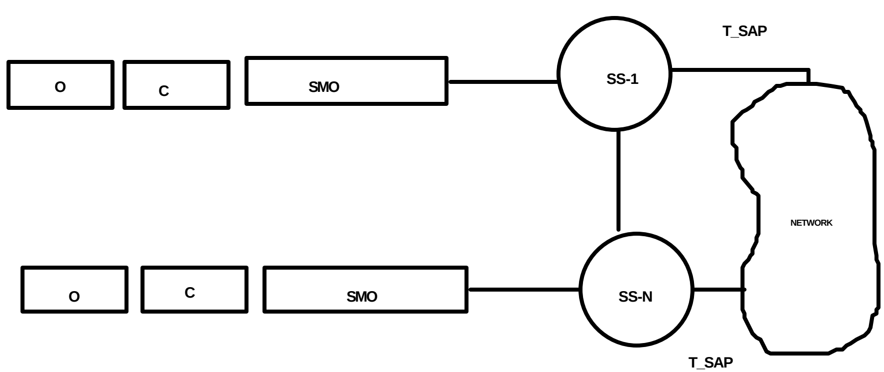

(ii) Multimedia Communications Protocols

The objectives of the communications system are thus focused on meeting the interaction between users who are communicating with CMOs. The implementation of these new communications systems will be driven by the implementation of new and innovative protocols. This section describes several of the driving factors that are necessary for their implementation. The protocols must be developed to specifically be able to perform the following tasks:

- o Allow any user to create an SMO and a CMO.

o Allow any user or set of users to share, store, or modify a CMO.

o Ensure that the user to user communications preserves the temporal, logical and spatial relationships between all CMOs at all users at all times.

o Provide an overall environment to define, manage and monitor the overall activity.

o Provide for an environment to monitor, manage and restore all services in the event of system failures or degradation.

The OSI layered communications architecture has evolved to manage and support the distributed communications environment across error prone communications channels. In general the layers 1 through 4 are directed at the transport channel or medium. A great deal of effort has been spent on developing and implementing protocols to support these channel requirements. Layer 7 provides for the applications interface and generally support such applications as file, mail and directory. The requirements of a multimedia environment are best met by focusing on layer 5, the session layer whose overall function is to ensure the end to end integrity of the applications that are being supported.

Some authors (See Couloris and Dollimore or Mullender) indicate that the session function is merely to support virtual connections between pairs of processes. Mullender specifically deals with the session function in the context of the inter process communications (IPC). In the context of the multimedia object requirements of the previous section, we can further extend the concept of the session service to provide for IPC functionality at the applications layer and specifically with regards to multimedia applications and their imbedded objects.

The services provided by the session layer fall into four categories:

o **Dialog Management:** This function provides all of the users with the ability to control, on a local basis as well as global basis, the overall interaction in the session. Specifically, dialog management determines the protocol of who talks when and how this control of talking is passed from one user to another.

o **Activity Management:** An activity can be defined as a sequence of events that may be within a session or may encompass several sessions. From the applications perspective, the application can define a sequence of events called an activity and the session service will ensure that it will monitor and report back if the activity is completed or if it has been aborted that such is the fact.

For example, in a medical application, we can define an activity called "diagnosis" and it may consist of a multiple set of session between several consulting physicians. We define a beginning of the activity when the patient arrives for the first visit and the end when the primary physician writes the diagnosis. The session service will be responsible for ensuring that all patients have diagnosis.

o **Synchronization**: We have seen that at the heart of a multimedia system is a multimedia data object. Each of the objects has its own synchronization or timing requirements and more importantly, a compound object has the orchestration problem. The session service of synchronization must then ensure that the end to end timing between users and objets is maintained throughout.

o **Event Management:** The monitoring of performance, isolation of problems, and restoration of service is a key element of the session service. Full end to end network management requires not only the management of transport and subnetwork, but requires that across all seven OSI layers, that overall end to and management be maintained .

The general architecture for the implementation of the session services (and all other OSI services) is shown in Figure 4.3. Here we have shown the session entity which is effectively a session service server. The entity is accesses from above by a Session_Service Access Point (S_SAP). The session entities communicate through a Protocol Data Unit (PDU) that is passed along from location to location. Logically the session server sits atop the transport server at each location.

The servers are conceptually at a level above the transport level. We typically view the transport servers as communicating distributed processes that are locally resident in each of the transmitting entities. This then begs the question as to where does one place the session servers. Are they local and fully distributed, can they be centralized, and if so what is their relationship to the Transport servers. Before answering these questions, let us first review how the session services are accessed and how they are communicated.

Session services are accessed by the higher layer protocols by invoking session service primitives. The model does not however say where the session server is nor even if it is a single centralized server, a shared distributed server, or a fully distributed server per entity design. From the perspective of the Internet, this begs the question of whether the session layer services as we have developed them are more appropriately done in the Internet as a Session Control Protocol, a SCP. Does this mean that the Internet must now deal with a SCP/TCP/IP layered system? The answer seems unequivocally to be yes!

4.2 Access Expansion

Access is defined as the physical connection between a local host or network entity and the body of the Internet. It is the last mile or the last 1,000 feet connection to the Internet. It also is generally the most costly. It may be the local telephone company and a dedicated access line, or it may be the campus local area network. We shall focus on the low end user exclusively, and shall consider access as a full time and re time connection to the Internet from a single low end user location to the remainder of the network. As a point of reference, PSI may charge a user almost $1,000 per month for such a connection, the costs dominated by the local telephone company. Such charges are the primary barrier to entry to the low end user. Access fees must be changed if the low end users is to be more than a sporadic users of the service This section discusses two alternatives that are currently being developed.

We shall focus on two of the most dramatic areas of change and discuss their architectural implications. Specifically CATV and wireless access.

4.2.1 CATV [8]

The existing CATV entities have argued that they have an infrastructure that is highly suitable to use for local access and broadband access. It has been previously demonstrated that use of this technology is viable in many areas but has it has also clearly demonstrated its severe drawbacks.[9] The current action

---

[8]The authors wish to note that when this paper was first published, May 5, 1993, there were no CATV access alternatives. In August of 1993, Continental Cable of Boston and PSI of Reston Virginia signed a joint agreement to allow access to the Continental Cambridge customer base to Internet at a proposed rate of $100 per month. This was a factor of 10 less than the telephone rates. As of this writing of this version of the paper there are no know customers on the proposed network.

[9]See the paper by McGarty and McGarty, 1983, describing the consumer trials with the QUBE system. The senior author was the first to develop and deliver a full motion video on demand system t customers over a cable and telco integrated system. In 1983 the senior author had a joint venture with Warner, GTE, Bank of America, Bell of Pennsylvania and DEC. This was a data over voice telco network integrated with the cable system in Pittsburgh. The system used a bank of 100 video disk players and was the predecessor of the current Warner 500 channel system in New York and Florida. The system used a space division multiple access scheme, splitting video packets on the head end-hub-sub hub networks.

between cable and telephone companies, namely Bell Atlantic and TCI, Cox and Southwestern Bell, and Time Warner and US West are clear indicators that there is some perceived joint interest. The argument has been for over ten years that cable can provide data and interactive services to the home more cost effectively than telephone. The argument has been operationally proven however.[10]

To demonstrate the potential of Cable for local access we review some of the Cable architectures and directions. The current CATV systems have a tree and branch architecture that is one way. We show this in Figure 4.4.

Figure 4.4 CATV Architecture

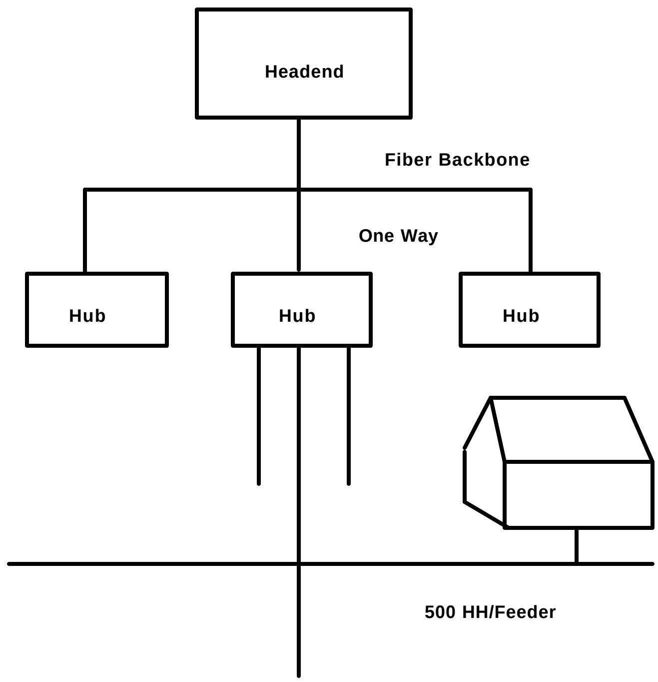

CATV systems have two major plant elements; the residential cable that services all homes and the institutional cable or loop which was required for th Cable franchises but never put to use. The institutional loops present significant access alternatives. CATV companies have also been adding additional backbone elements and have been using these as CAP, Competitive Access Providers, in the commercial areas.

CATV was initial two way capable and some systems such as the Warner QUBE systems were two way activated. However, this requires the use of two way repeaters which are inactive in almost all plants. There is 50 MHz of bandwidth available in all CATV systems for residential use and 150 MHz for all institutional use. The question is how to take a one way system and make it two way. This is shown in Figure 4.5.

[10]The classic case is the example of Manhattan Cable, now Time Warner cable in New York City. It was the first such entity to deliver 56 Kbps service. It was out marketed by Teleport and by MFS. It had an early strategic position, but due to its lack of understanding of the telephone business, and more fundamentally, the fact that cable is not a telecommunications system but and entertainment distribution system, its availability and reliability was too low for commercial applications. It has since been reduced to a non-player in the market.

Figure 4.5 CATV Two Way Plant

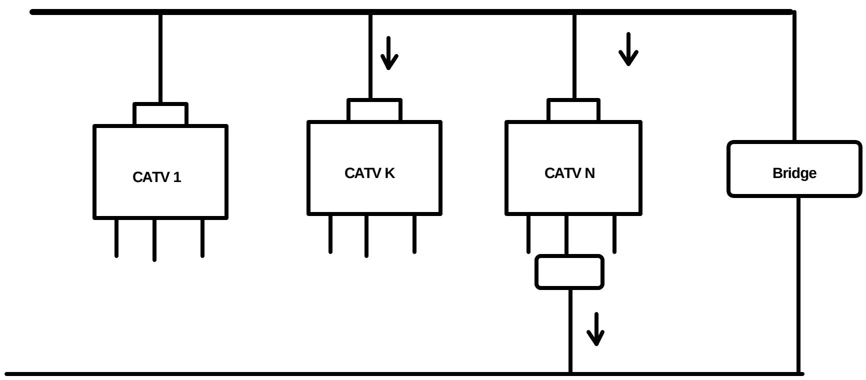

This two way capability is achieved by creating a 50 Mbps or even 100 Mbps bus comparable to that provided in FDDI.[11] This is achieved by looping the local feeder cables back to a bus fiber via a bridge. Thus the system looks like a broadband system allowing any home or location with up to 100 Mbps service with no capital change in the CATV plant. The question then is, what are the user requirements for such a system. Simply put, we can assume that a single user is accessing the system and that when accessing it transfers 4Mbit of files per hour. This is approximately 1Kbps average rate, and on average if there are 8,000 users, we need about 24 Mbps to handle all of them.[12] Thus a cable system can theoretically handle a reasonable community of users.

In a CATV environment there are several reasons, however, for lack of CATV infrastructure at the present time:[13]

(1) **Interconnect**: In a reasonable radius from any large metropolitan area there are one to several dozen CATV entities. The issue of interface and interconnect has never been adequately addressed and there are no standards that allow for this. In addition, CATV switch access uses the same dated architecture as does cellular and thus is highly reliant upon the existing LEC. This will merely drive up the costs of goods for the carrier.

(2) **Availability**: CATV systems have system availability numbers that are less than 90%, whereas communications networks have availability numbers in excess of 99.5%. The inherent structure, operations and management of the two networks are currently incompatible. Specifically CATV, as currently operated cannot provide toll grade quality service.

(3) **Bandwidth**: Bandwidth in a CATV system is limited, except on Institutional loops. Local bandwidth is structured for video and the two way systems have limited return path. When we begin the multimedia applications, the bandwidth is even greater.[14]

---

[11]Cable systems are generally split with a 50 MHz band on the return pat, the path from the home to the head end. All of the other bandwidth is from the headed to the home, see McGarty and McGarty, 1983.

[12]See McGarty, Access Policy, Telecommunications Policy Research Conference, October, 1993, Solomon's, Island.

[13] McGarty, T.P., R. Veith, Hybrid Cable and Telephone Networks, IEEE Comp Con, 1983. and, McGarty, T.P., S.J. McGarty, Impacts of Consumer Demands on CATV Local Loop Communications, IEEE ICC, 1983.

(4) **Performance**: Data transmission performance on coaxial or fiber/co-ax has been shown to have significant problems due to an excessively noisy environment resulting from many open cable access termination in homes of current or prior subscribers. Admittedly this may be ameliorated but it will require significant rebuilds as well as management and administration of the subscriber loop.

(5) **Inactivated Two Way Returns**: Two way cable almost ceased to exist as an operating entity with the demise of the famous QUBE system. [15] Currently there are less than 0.1% of the CATV systems with active and operational cable return paths and supported bi-directional amplifiers.[16] For the CATV system to function this must be addressed.

However, we argue that these limitations can be eliminated from the current designs. Specifically,

1. *Interconnect*: CATV companies are being pressured to establish interconnect amongst themselves at the head ends as well as entering into the CAP, or Competitive Access Provider, business. This will ensure system to system access. Internet access can be achieved directly via a CAP point of presence, or PoP.

2. *Availability*: Trunk CATV transport is being upgraded to fiber which is much more reliable. By establishing a bus network on top of that, alternative routing is possible and the system availability improved.

3. *Bandwidth*: Using the Institution loops, 150 MHz is available and with high signal to noise and controlled interference it is possible to achieve 300 Mbps on such loops. Thus a bus of 300 Mbps or more is achievable in today's architecture.

4. *Performance*: This is a major problem area generally due to open taps on the cable. This may require better maintenance and control, which is a methods and procedures problem for the CATV operators.

5. *Two Way*: This is achieved via the bus proposal.

Therefore, CATV provides Internet with a viable and current option for the expansion into high data rates and multimedia capability.

4.2.2 Wireless

Wireless communications services introduce new sets of technologies that will create a new local loop access paradigm. The current view of the local loop is that of a bundled set of services that possess significant economies of scale and thus justify permitting the Local Exchange Carriers (LECs) to have a total monopoly in the local exchange. The new technologies allow dramatically lower capital costs per subscriber and also eliminate the scale and scope economies in local access. The only remaining stumbling block is the

---

[14]McGarty and Clancy. The author develop the model for integrated video, voice, text ad data on an integrated cable-telco system. This paper, although for proprietary reasons not describing the full motion video on demand system, does present the detailed technical analysis.

[15] McGarty, T.P., G.J. Clancy, Cable Based Metro Area Networks, IEEE Jour on Sel Areas in Comm, Vol 1, No 5, pp 816-831, Nov 1983. and McGarty, T.P., Local Area Wideband Data Communications Networks, EASCON, 1982.

[16]For example, the author has been told by Continental that Cambridge was the only system with two way activated feed in their entire network.

cost of switch access, called the access fee. This access fee currently represents an expense equal to fifty percent of the gross revenues of AT&T, MCI and the other IECs. Access fees account for one third of the gross revenues of the LECs. These fees are the last remaining barrier to entry in providing fully competitive communications services.

We discuss the issues concerning local network access from both a technology and policy perspective. The technology shift is discussed showing how the current copper wire based network having high capital per subscriber values is being replaced by a wireless network having capital per subscriber numbers almost two orders of magnitude smaller. [17] The basic microeconomic implication is that there is a loss of economies of scale and scope that were at the heart of the old "Bell System." The technology shift demands a policy shift on the part of regulators and government officials.

The differing policy options are discussed and their implications on the telecommunications industry evolution are presented. The concept of disaggregation of the LECs into wholesale switch and transport companies is presented in the context of the current policy discussions on PCS at the FCC. The imp act on new markets and services as a result of eliminating the local bottleneck is discussed . Key policy issues relating to providing a seamless interoperable national wireless network are presented. PCS, via a dramatic technology change, will create new business opportunities and will also challenge the policy makers and regulators to move quickly but carefully in developing and implementing these new directions for telecommunications evolution.

Technology has changed dramatically in the past five years. The two current ways of providing voice service are via wireline twisted pair telephone service and through cellular voice service.[18] New technological innovations have allowed the wireless PCS services to be provided by another form of technology. This technology takes advantage of a distributed telecommunications architecture and places as much "silicon" in the field as possible. It also performs as much processing as possible so as to minimize the functions required by the LEC interconnect.

We shall use the example of CDMA technology to demonstrate how this new technological infrastructure can enable the new market. We shall briefly describe the CDMA system and then proceed to the financial implications of using this new technology. The CDMA system described is that of QUALCOMM [19]. Fundamentally the system is characterized in the following fashion, and as shown in Figure 4.6:

(i) An air interface of a CDMA signal is provided by a cell or cell re-rad over the air to the portable. The signal is encoded in a direct sequence CDMA spread spectrum code. Thus a 9.6 Kbps signal is spread, or multiplied by a unique code at the rate of 1.25 Mbps. The codes are orthogonal. Namely, if two or more codes are combined, then if they are multiplied by the desired code, the residual of the other signal appears as a low level noise signal. Thus CDMA is frequently interference limited no random noise limited.

(ii) A cell controller is used to ensure hand off between other cell controllers. The cell controller has a capacity that depends upon the bandwidth, the interference level, the size of the cell and other factors. Typically a cell controller has the capacity of 500 to 1,000 trunks. Note that this is given in trunks and not portables. If a portable is busy 5% of the time then this means a cell

[17] de Sola Pool, 1977 and 1990. The 1977 text gives a detailed description of the history and structure of the telephone market.

[18] See the works by Lee. The author has provided several key bodies of analysis that provide insight into the history and current status of cellular.

[19] See the works by Gilhousen for the QUALCOMM approach. Also see the paper by Pickholtz et al for a differing approach to CDMA. The latter approach is Broadband CDMA compared to mid-band.

controller with 1,000 trunks can handle 20,000 portables. The cell controller is a highly intelligent distributed processing node. The CDMA codes assure signal orthogonality and inherently manage the interference. The cell controller assures a soft hand-off between the other cells in the grid. In addition, the cell controller establishes the relationship between the call and the switch. Namely the cell controller passes an intelligent and digitally "packed" set of voice channels.

(iii) The cell-controller hands the switch a DS-3 formatted voice signal, with a SS-7 signaling channel, on a SONET interface. As far as the switch is concerned, the call may have originated from a Class 5 or Class 4 switch. As we have discussed before, the Class 5 LEC functionality is not required. What is required is the Class 4 toll-tandem switching capability. The only need for Class 5 functionality is that needed for billing.

(iv) The re-rads are clustered around the cell controller. A re rad is used to manage the coverage issue, whereas the cell controlled is used for the capacity issue. The re-rads are an order of magnitude less expensive than the cell controller. The re-rads are interconnected to the cell controller via a microwave path, at 40 GHz, or over CATV or a bypass carrier.

Figure 4.6 Wireless Architecture

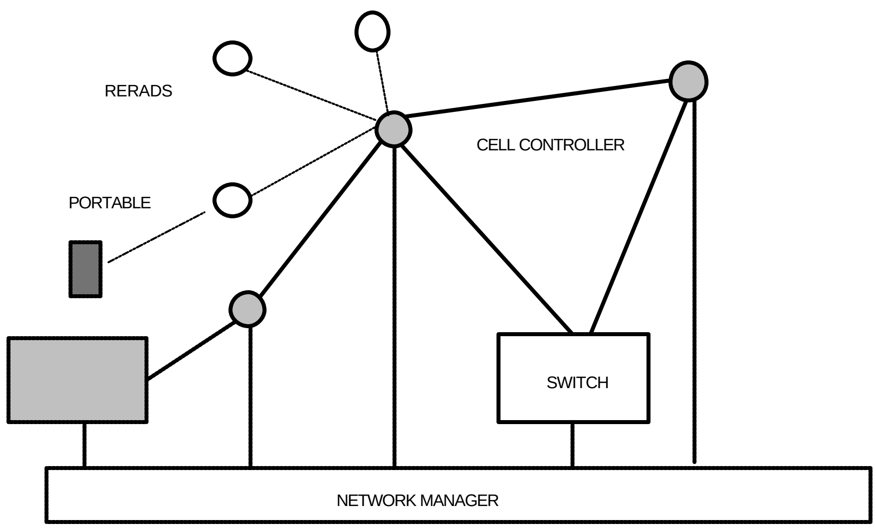

When looked at in this fashion, the use of CDMA dramatically reduces the needs from a LEC environment. All that is needed is the ability to backward access to the local user, namely a customer of the LEC. Thus the access fee should be reduced.

A simple calculation will show how this new technology dramatically reduces the capital per subscriber.

o Assume that there are 1,000 square miles of coverage and 48,000 subscribers.

o Assume that a cell controller or a re-rad handles a 3 mi. radius or about a 30 mi. cell coverage area. This implies that 3 cell controllers and 30 re-rads will cover the area.

o Assume that the cell-controller is equipped to handle 800 trunks per cell controller. Assume that the peak usage ratio is 5%. Thus each cell controller can handle 16,000 subscribers, 800 instantaneously active in the busy hour.

o Assume that the cell controller are about $1 million each and that the re-rads, with microwave back haul are at $50 thousand each. The total capital is $4.5 million. This the is about $100 of capital per subscriber.

Now this can be compared to the capital per subscriber in the LEC and cellular environments. In the LEC world the capital per subscriber is almost $1,800. This is split between switch and transport as follows; $500 for the switch and $1,300 for transport. Namely, the LEC is outside plant dominated. Moreover, under rate of return regulation, the LEC makes most of its profit off of its outside plant. In the cellular world the capital per subscriber is $750. This includes the cells and the MTSO, Mobile Telephone Switching Office. It does not include access to the LEC Class 5 switch.

What does wireless do for the Internet? There are several things that it does immediately:

Access Expansion: Providing access to places not readily served. This is, for example, school class rooms and other locations not now served by wire based telephone.

Expanded Bandwidth: Using wireless, it may be possible to take 20 to 40 MHz of bandwidth, and create a 40 to 80 Mbps bus to allow PDA users access to a wide variety of services including multimedia.

Terminal Identification: Wireless has an infrastructure that will enable the "Find Me!" paradigm to be effected. It has more than the cellular roaming capabilities that we have seen evolve in the older analog cellular architecture.

It is clear that wireless will introduce many new dimensions to access.

4.3 Host Migration

Host Migration is the term used for the identification of user with host and vice versa. The changes occurring in host performance and characterization that forces consideration of this changes and its implications are as follows:

Higher Performance Single Processors: The development of higher speed and performance workstations has dramatically changed the capabilities of the local host. The individuals work station will have full host capability and interconnectivity and the identification of the host will migrate down to the lowest common denominator, although those machines will have significant power.

Personal Digital Assistants: The PDAs will create a dramatic change on the Internet. The PDA is a wireless device that identifies with the individual. It is as powerful as many other host like processors but it dramatically changes the paradigm. It associates itself with a person, at all times, and not with a location. The difference is between that of the existing telephone network and the new PCS telecommunications network. The PDA will challenge the Internet to recognize the use qua user rather than the user qua host.

Wireless Host Access: Wireless access presents challenges of expansion, connectivity, interconnectivity, and addressability. In 1975 and 1976, the senior author was responsible, with others, for establishing the first wireless access to the ARA Net via Intelsat. The challenge at that time was connecting the earth stations in West Virginia to that at Goonhilly in the UK. This was one of the first wireless connections. Earlier Kahn and Abramson had developed the Packet Radio approach but had limited ARPA Net Access. The challenge in both of these cases was the need for the development of new data protocols because of the unique capabilities of the channel. We shall see the same here but now we must expand to SCP/TCP/IP.

Multiple Host per Single Users: The user will have the identity and the hosts may be connected to multiple users. Moreover there may be a set of many to many connections. The issue of "Find Me!" will be expanded to connect to and communicate with the user "Now!".

Enhanced Host Access Software: The end user currently access the Internet via the UNIX like calls and commands. A few simplified front end devices are available but are limited. The expansion of low end user host access capabilities, such as through a Windows type environment, albeit in a multitasking mode, will dramatically increase the demand. Add to that a CATV or wireless packet access mode rather than a twisted pair local lop, tariffed at a per packet or even fixed rate basis, will dramatically expend the host migration.

5.0 Microeconomic Implications

The access of the low end users raises the question of costs for the service and the resulting pricing. This sections presents a brief micro-economic analysis of the Internet access

5.1 Cost Structures and Cost Allocations

There have been several attempts to analyze the cost structure of the Internet, such as that by Fulhaber and that of MacKie-Mason and Varian. However, none of these studies have truly looked at the Internet from a micro perspective, one that would be used in the development of a business or operations model of the business. In this section we develop the structure for such a model but are not able at this time to adequately ascribe the detail necessary to determine the value of all of the parameters presented. This is a process driven model of the business.[20]

We shall develop the model using the functional approach and from that use the productivity, market driver, and unit cost model. We divide the system into Expense elements and Capital elements. Let us first begin with the Expense Elements. In so doing, we shall introduce expense elements that may not now currently be part of the bundled Internet costs, but maybe concomitant with migration to a low end user community.

The Expense Elements that should be included in the detailed costs analysis of the Internet are as follows:

> **Engineering**: This includes all of the allocateable costs and actions relating to the engineering of the system, not including the implementation of the elements. The engineering functions include all such tasks as product development, service development, standards setting and other such tasks that result in the delivery of a detailed description of the tasks to be performed.
>
> **Development**: This includes all of the tasks related to the actual implementation of enhancements, improvements, and alterations to the network.
>
> **Installation**: These are all tasks related to the installation of the network capital plant and the network software. They include all actions taken after the completion and acceptance of any beta prototypes.
>
> **Operations**: This includes all ongoing tasks and efforts related to the operating of the network.
>
> **Maintenance**: These are all tasks that relate to the correction of faults or failures on the network. The efforts relate to how well the network functions as measured by such factors as mean time to failure, mean time to repair, and other comparable factors.

[20]The model is based upon the methodology developed by McGarty, Business Plans, and used extensively in modeling communications and information system Architectures as well as operational systems.

Customer Service: This is the set of all functions related to the support of the customer base. It can be related to the number of calls per customer per unit time and the holding time per call.

Billing: These are all of the tasks allocated to the process of billing and collecting of the revenue. If there is a fixed price billing then this is simple. If there is a usage sensitive billing then this could be a complex operation.

Marketing: This is the collection of all non technical activities related to the development and specification of new services.

Sales: These are all direct and indirect sales activities.

Administrative: These are all allocateable administrative costs of all functions not in the above defined tasks.

General Management: These are the costs of General Management not in any of the above.

Overhead: These are all expenses related to or ancillary to the operations such as rent, insurance, travel etc.

We develop the model in the following fashion. For each expense area, K, we define a fixed and variable amount. The fixed amount is not considered in the analysis directly. The variable amount is driven by what we term a revenue driver, RD. The revenue driver may be the number of customers, the number of customers per segment of market, or the number of networks, the number of hosts or the combination of these.

The second parameter is the productivity factor. If we consider customer service then the driver is the number of customers. The productivity factor is a factor of the calls per year per customer and the holding time per call. This will give an Erlang load of the customer service function, or an average number of hours per week, per thousand customers. From the perspective of an employee of 40 hours per week, then we find out how many customer service personnel we need per thousand customers. This is the productivity factor.

Finally is the unit costs or the salary the customer service person. The model follows as below:

$$E_k = \sum_{n=1}^{N} RD_{k,n} PF_{k,n} UC_{k,n} + FC_k$$

Furthermore, the total expenses can be written as:

$$E = \sum_{k=1}^{K} E_k = \sum_{k=1}^{K} [\sum_{n=1}^{N} RD_{k,n} PF_{k,n} UC_{k,n} + FC_k]$$

Now we can relate this to the Expense parameterized on the number of users, n, where we have the following:

$$E(n) = \sum_{k=1}^{K} E_k(n) = \sum_{k=1}^{K} [\sum_{n=1}^{N} RD_{k,n}(n) PF_{k,n}(n) UC_{k,n}(n) + FC_k(n)]$$

It can be shown that the total expense can be modeled in this fashion. We will not perform the analysis her since most of the data is not yet available.

In a similar fashion we can do the same for the capital structure. If we define the capital elements in terms of such items as routers, switches, fiber, access nodes, etc., and then develop market drivers such as the number of users, number of hosts, number of users per host, etc., and then use performance factors such as the number of users per router, and finally the capital per router. Thus the capital follows the same approach, but now:

$$C(n) = \sum_{k=1}^{K} C_k(n) = \sum_{k=1}^{K} [\sum_{n=1}^{N} RD_{k,n}(n) PF_{k,n}(n) UC_{k,n}(n) + FC_k(n)]$$

Which has a form comparable to the expense item.

5.2 Scale and Scope

We can now address the issues of scale and scope in this business. Scale implies a reducing marginal costs and average cost as we increase the number of users. Scope implies that there are exogenous factors in some other business element that cause reduction in cost per unit. We shall not focus on scope suffice it to say that it is an issue that will require additional analysis in the context of the Internet as an entity in a broader commu nity.

As to scale, let us define the average cost as:

$$AC = \frac{E(n) + C(n)}{n}$$

and the marginal cost is:

$$MC = \frac{\partial(E(n) + C(n))}{\partial n}$$

where C is the capital at unit level n.

Also recall that the supply curve for this business is the marginal cost curve above the point of minimum average cost.[21]

Let us take the simple example for the Internet of the customer service function combined with a fiber backbone. Customer service expenses are driven by the number of users and the traffic (or Erlang) load per customer divided by the hours per staff person, all multiplied by the cost per service person. Specifically:[22]

---

[21]Pindyck and Rubinfield, p. 254.

[22]This states the following: The arrival rate is the average number of calls per hour to the customer service center from the customer base. The holding time is the amount of time it takes to handle a call on average. The product of average call rate and average time per call is a measure of the calling traffic lo ad, usually expressed in Erlangs.

$$E_{CustomerService} = n\left(\frac{ArrivalRate * HoldTime}{StaffHours}\right)Salary / Staff$$

Thus the averages and marginal costs are the same, namely;

$$AC = MC = \left(\frac{ArrivalRate * HoldTime}{StaffHours}\right)Salary / Staff$$

if the arrival rates and holding times are independent of the user mix. However, we know that as we expand the base the number of inquiries increase and the length of time of the inquiries also increase. Thus the AC and MC are not the same. In addition, there may be some allocate fixed costs. The AC can be given by:

$$AC = AR(n)HT(n)\left(Salary/Staff\right) + \frac{Fixed}{n}$$

The first term is the average variable cost. Let us examine its dynamics. At the low end of penetration, n being small, the arrival rate, AR, and the holding time, HT, may be large because the service in new and there frequently are many problems understanding the service. At the high end with many uneducated customers, the AR and HT may again be large. In the middle, the average customers have been slightly educated and moderately self dependent, and the service mature and stable. At that point the VAC, variable average cost, is minimized.

The capital for fiber is a similar analysis, but it moves in steps. A fiber strand may be placed in and it has a dramatically fixed capital that makes the average continue to decrease and yet be far from the marginal. Thus the fiber part of the equation is dramatically different in concept but similar in result.

5.3 Supply and Demand

As we stated, the supply curve can be determined from the above analysis, such an analysis yet to be done in the reported literature. We now want to combine the supply and demand analysis to show how, if this data were available, a more detailed analysis for policy objectives could be determined for the Internet.

The demand curve has several components. The high end component is an indirectly subsidized entity, since the users are paying via their institutions or via their grants, and not necessarily directly and personally. We have taken liberties with the term subsidy here, to establish a distinction between a commercial or institutional user and a low end direct user. One of the distinctions is that of cost or price proximity. For the most part the high end user does not see the payment.

The low end user sees the payment directly and the payment policy will influence the penetration in this market more directly. The economic model for the Internet is displayed in Figure 5.1. We argue that it is a monopoly priced service and show this in terms of the equilibrium point of marginal cost, MC, equals marginal revenue, MR.[23] This differs from a totally competitive market facing a fixed price point.[24]

---

[23]We also argue that the MC is upward sloping, which may or may not be the case depending on the equilibrium point of the market.

[24]As is shown in Pindyck and Rubinfeld or Henderson and Quandt, Profit is given by the following:

$$\Pi = R - C = pq - C(q)$$

Figure 5.1 Internet Economics; The Quasi Monopoly Network

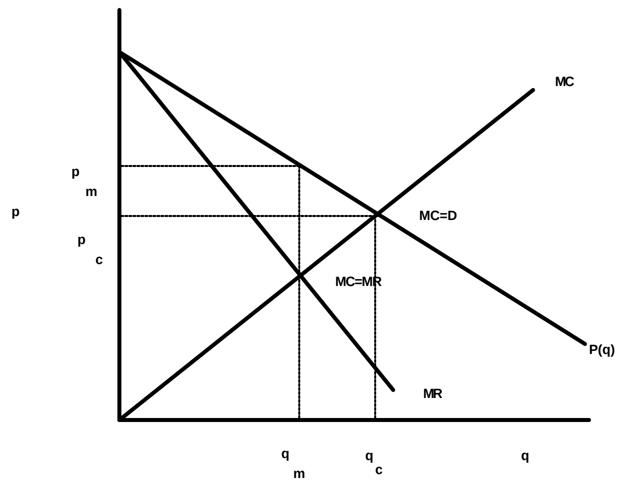

The question that must be asked is; the Internet is really a monopoly or is it some other form of market entity? Clearly it is not a fully competitive player in the market. One cannot name another Internet type system of similar functions and size. However, it is not truly a monopoly, not being a single economic entity controlled by a single set of management. In contrast, it has been argued by Einhorn that the Internet resembles the Gas and Electric Utilities with their backbone networks.[25] In fact, as Einhorn notes, this analogy has many significant insights for how the Internet may evolve as the backbone portion of a more distributed and competitive network of information interconnects.

Now we know that the Internet has an effective subsidy, s, that currently is given by the Government. We show this effect in Figure 5.2. The subsidy increases demand, reduces the consumer prices and subsidizes the supplier, in this case Internet, as an entity. The subsidy clearly benefits all of the user base. The question is; can the demand be maintained as the subsidy is eliminated? Also, can the subsidy be measured from the perspective of the total user base, namely there are impact subsidy elements that are not directly reflected, namely, overhead costs supported on Government Research Contracts.

---

and that profit is maximized by the value of q that meets the following condition;

$$\frac{\partial \Pi}{\partial q} = 0 = \frac{\partial (pq)}{\partial q} + \frac{\partial C(q)}{\partial q}$$

or the MC.

[25]Einhorn, Michael, Personal Communications.

Figure 5.2 Internet Subsidy

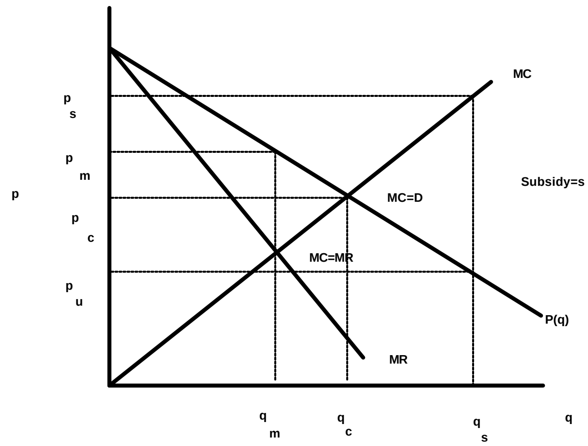

Finally, we can view Internet market dynamics from the perspective of two issues. First we see what happens to Internet when we change the access fees. This is a change in the supply curve. Second we change Internet with the addition of Multimedia. This changes both the supply and the demand curve. The access fee changes may have little effect on the higher end users who have not seen them in detail of more frequently have them bundled in a more cost effective fashion.

On the other hand, for the low end user, the access fee, say of $1,000 per month, if reduced to $50 per month, could result in dramatic demand shifts. For example, a small company called Air Access has developed a wireless large area Ethernet Internet access that is tentatively priced at $50 per month. The demand projections for this service are dramatic. This type of access costs will create the low end user access. It may be an added convenience for the higher end user, but may not change their usage dramatically.

Multimedia, in contrast to access which shifts the supply curve, shifts the demand curve. It has been estimated that the addition of true multimedia communications as described in this paper could readily double demand for the low end user market for the same costs of service price point. We show these in Figures 5.3 and 5.4. The issue to note is that there can be a significant shift in overall usage. It will beg the question of continuing sustained subsidy.

Figure 5.3 Reduction of Access fees on Internet

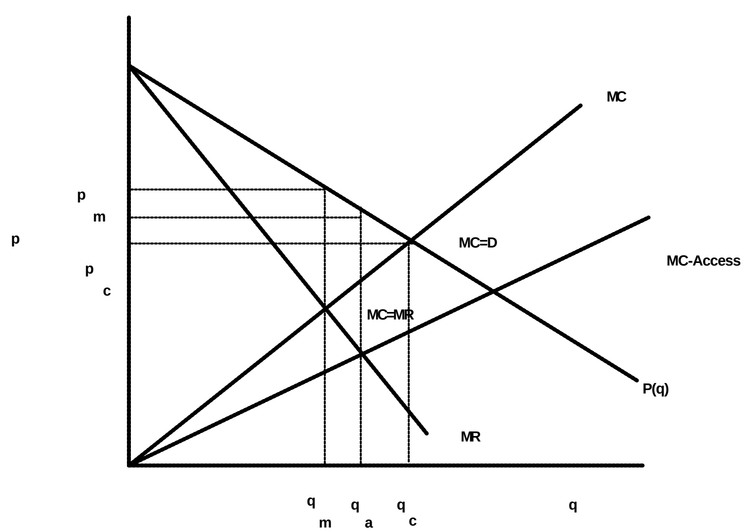

Figure 5.4 Introduction of Multimedia on Internet

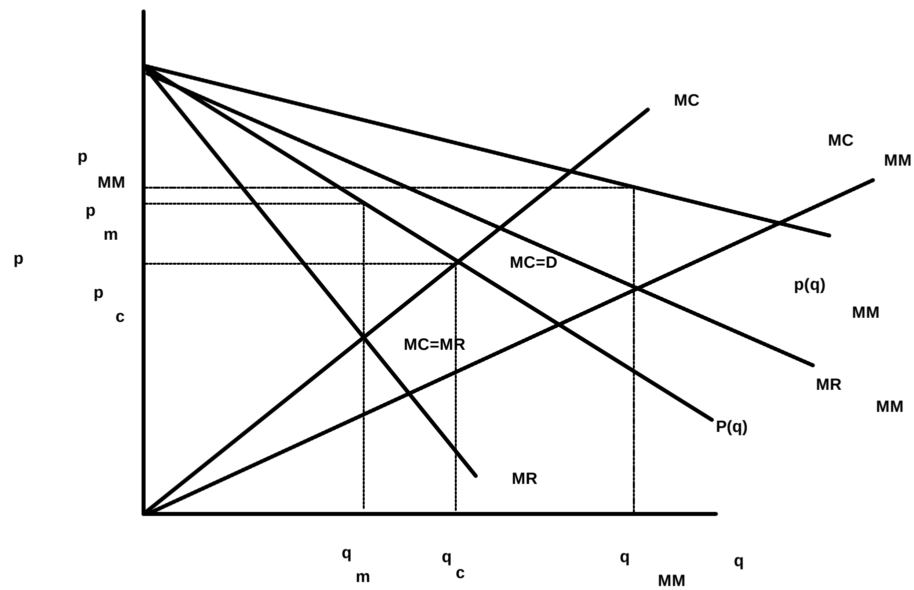

In Figure 5.4 we show the impact of new services, thus increasing the demand curve, combined with new access, thus decreasing the supply curve. The market and technology dynamics of the Internet require more intense study to better understand these overall microeconomic dynamics. The analysis presented above, clearly requires significant detail to be through.

Yet it does raise the structural base for the detailed analysis of the key policy issues. Figure 5.4 depicts the interaction of cost reduction and service enhancement. It begs the question of the cost of the service enhancement. We have argued elsewhere that such costs are reflected more in the end user terminals than in the Internet infrastructure.[26] Specifically, multimedia communications is software resident and not silicon resident. Silicon is almost "free" is the main driver for the continued costs reduction. The software costs must, however, be accounted for. Yet even there, it can be argued that with lowered access costs, and distribute multimedia communications in a client server design, the user will see an increased demand and lowered cost.

6.0 Policy Implications

There are several areas that the low end user impacts upon in terms of policy decisions. Specifically:

Access:

The issue of access is an issue of flexibility and ubiquity versus an issue of regulation and revenue assignment. The concern is that with multiple means to access the Internet for the Low end user, and the high end user as well, the access fees that are currently charged by the LECs reflect a world view that is of a pre divestiture nature and attempt to continue to support the local user at the expense of the long distance user. With the Internet being the long distance company, in effect a pseudo IEC, and the local access being separate from the LEC, does the user have a requirement of continuing to reimburse the LEC for access at the other end independent of who placed the call. The answer should be no. The issue is, however, not that

[26]McGarty, IEEE, November, 1992.

simple. The argument is changed significantly when the alternative access schemes are introduced. If one uses a wireless scheme or a cable scheme, access is eliminated and is subjugated to competitive market pricing and not monopolistic Local exchange Carrier pricing.

Pricing:

There is the issue of pricing on a value base versus a cost base. As we have discussed in the earlier sections, information has an intrinsic value that can be measured in a variety of ways. The second policy issue of pricing is that of fixed pricing versus variable pricing. The third pricing issue is that as relates to the settlements process in the access area. Namely, if there are multiple access regimes why should there be settlements if the traffic is "generally" balanced. The balancing of traffic is a key issue that needs addition work. It begs the question of information sources and sinks. For example, if traffic were not balanced, namely that on average input equals output, there would conceptually be a net accumulation of traffic, read information, at some preferred nodes on the network.[27] Finally, from a pricing perspective, there is the question as relates tot he low end user of Universal Service and a pricing schedule that support that. This begs the question of subsidies and "taxes" .

There are two Pricing Regimes that are of some concern. The first is the "Limited Capacity" regime wherein the system is backed up to its limits and there is a premium placed on any one users demanding increased capabilities. The second regime is the "Unlimited Capacity" regime wherein the backbone, the regionals and the campus nets are all over designed due to some artifact. In the latter case we will argue for fixed pricing and unlimited Class usage. In the former case we argue for Premium based pricing.

Universal Service:

As we stated in the Pricing Policy area, there is a Universal service policy question. Who gets the service, who pays, and at what prices? This is the case of K-12 and other areas of users access. Should there be a subsidy to prime the market pump. Who should pay that subsidy. The argument made by the authors is one that states that universal access is a social policy. The network entities may be required to implement a social policy but the costs of implementation should be paid via a taxation mechanism executed by the Government and not by the network entity. The social policy should be clear and articulated and agreed to as such, the taxation should also be so stated. The any other network entity should have the same responsibility.

Security and Privacy:

The Internet is a Global Network and as is projected there will be a one to one marrying of host and user. This begins to beg the question about security in this environment. If the users can now be acccesd at any time and place, via a wireless capability and into a PDA, and thus networked to the users other hosts, what security access controls are there to be in the network. As the SCP or Session capabilities place more intelligence in the network, including the backbone, what safeguards must be developed to ensure integrity and privacy. It is argued that these new movements into multimedia, access and host migration engender new and different concerns about security and privacy.

---

[27]McGarty, Access Policy, TPRC, 1993. The author has argued that traffic in any communications system that is spread across a generally mixed user base, with minimal node biases, has traffic that allows the network to be cut at any link, or sets of links, and that the traffic into and out of the links, is on average, the same. The author then uses this argument, based on actual traffic statistics, to argue for the elimination of all access fees, for comparable local exchange entities. We argue the same here with a similar extension of the traffic construct.

Nodal Control:

Nodal control is the rephrasing of the issue of end user migration and the movement of client server functionalities into higher layers of the Internet architecture. The nodal control issue is where does the server capabilities reside, who controls them and how does all of the other policy implications relate to their use by the low end user community.

Carriage Issues:

The carriage issues raise several questions. The first is what type of Carriage is the Internet? Is the Internet ultimately a common carrier? If it is, what does that imply about access and access control and discipline. We argue that the Internet is truly a common carrier as stated in common law. As a common carrier the Internet has the responsibility of being an open network with open interfaces. This is readily achieved by the workings of the Internet Society which sets the standards and establishes the overall architectural evolution. The Internet in unique in its ability to deal with users in a fully distributed format.[28]

Authorized Use and Users:

Authorized users and usages are coupled issues. As the Internet is opened up, as users are directly connected through host migration, and as access is commonly accepted as a real time and ongoing process, the question is raised as to who is a user and what are the appropriate uses. Is there to be a distinction between low end users and commercial user. Is there s pricing differential. At what point is a fixed access fee appropriate and when is it required to change.

These are several of the emerging policy issues that relate to this architectural migration of the Internet into the low end user community.

7.0 Conclusions

This paper is a first attempt to layout and describe the options of expanding the Internet into the low end user community. It stresses that issues of market demand and revenue potential, and then talks about the uses of the network to this community. We then address the issues of architectural changes and stress the concern of protocol "creep" in the network to the higher layers. To handle the new low end user community it will be necessary to deal with higher layers in the Internet, layers that have not been focused on in prior generations. We then discussed the economic structure of the Internet. It is clear that there are no detailed cost models for the Internet and that this is a serious lack that must be brought together. The macroeconomics of the Internet are then discussed. Lacking a strong economic model for the Internet at best we cane generate broad generalizations of issues. Finally we discuss the set of policy implications that result. There are many such issues that must be addressed. We must deal with them to effectively migrate the Internet into this new community.

The key points raised in this paper fall into the following categories:

- *Technological Innovation and Architectural Change: There are major technological changes that will allow increases in the high end user base but more importantly will empower the low end user. The changes will redefine the user and will let the user become a more integral part of the Internet.*

[28]The Internet is a network wherein the organization that manages the network reflects the form of the network. The organization is an open distributed form unlike the hierarchical form of management of the RBOCs, which in turn have a hierarchical network.

- *Structure, Pricing, and Market Behavior: As the user base migrates and moves into segments wherein the value of the information transacted across the Internet varies, the issue of fair and equitable pricing becomes a key issue. To price we must know costs. To know costs we must understand the micro structure of the Internet and its operations. There is at present no clear understanding of that issue.*

- *Policy Evolution and Alternatives: Policy development in the context of the Internet must be phrased in terms of objectives, such as low end user expansion, and in recognition of the architectural evolution. The issues raised in this paper as regards to policy are but a few of the many issues relating to the long term policy implications of the Internet.*

Acknowledgments:

The authors would like to thank several individuals who helped in discussions leading to this paper. Vint Cerf, of CNRI and The Internet Society, has provided many insights into the Internet, both history and evolution. Rich Mandlebaum, of the CATT Center at Polytechnic University, has provided the perspective of the network and Regional player as well as insight on the user community. Gad Selig, of ANS, has added the perspectives of ANS and their implications as the Internet migrates a more commercial segment. Richard Solomon, of MIT, has provided insight into pricing methodologies and mechanisms. Brian Kahin, of the Kennedy School at Harvard, continued to provided the insight into many of the policy areas. David Tennenhouse and David Clark, both of LCS at MIT, have provided a sense of direction of the TCP/IP migration to handle the challenges of multimedia.

References

Brenner, D.L., Law and Regulation of Common Carriers in the Communications Industry, Westview (Boulder, CO), 1992.

Cerf, V., Internetworking and the Internet, CNRI, 1993.

Christodoulakis, S., et al, Design and Performance Considerations for an Optical Disk Based Multimedia Object Server, IEEE Computer, Dec 1989, pp 45-56.

Dertouzos, M.L., J. Moses, The Computer Age, MIT Press (Cambridge, MA), 1979.

Fulhaber, G., Pricing Internet: The Efficient Subsidy, in Building Information Infrastructures, B. Kahin Ed, McGraw Hill (New York), 1992.

Gilder, G., Telecosm, Forbes, ASAP, March, 1993.

Hanson, Owen, Design of Computer Data Files, Computer Science Press (Rockville, MD) 1988.

Henderson, J.M., R.E. Quandt, Microeconomic Theory, McGraw Hill (New York), 1980.

Huber, P.W., The Geodesic Network, U.S. Department of Justice, Washington, DC, January, 1987.

Kahin, B., Building Information Infrastructure, McGraw Hill (New York), 1992.

Kahn, A.E., The Economics of Regulation, MIT Press (Cambridge, MA), 1989.

Krol, E., The Whole Internet, O'Reilly & Associates Inc (Sebastopol, CA), 1993.

Lehman, D.E., D.L. Weisman, Access Charges for Private Networks, Interconnecting with Public Systems, Twentieth Telecommunication Policy Research Conference, September, 1992.

Little, T.D.C., A. Ghafoor, Synchronization and Storage Models for Multimedia Objects, IEEE Journal on Sel Areas in Comm, April, 1990, pp. 413-427.

Lynch, D.C., M.T. Rose, Internet System Handbook, Addison Wesley (Reading, MA), 1993.

MacKie-Mason, J.K., H.R. Varian, Some Economics of the Internet, Tenth Michigan Public Utility Conference, March, 1993.

McGarty, T.P., Local Area Wideband Data Communications Networks, EASCON, 1982.

McGarty, T.P., S.J. McGarty, Impacts of Consumer Demands on CATV Local Loop Communications, IEEE ICC, 1983.

McGarty, T.P., G.J. Clancy, Cable Based Metro Area Networks, IEEE Journal on Selected Areas in Communications (JSAC), Vol 1, No 5, pp 816-831, Nov 1983.

McGarty, T.P., Business Plans, J. Wiley (New York), 1989.

McGarty, T.P., Alternative Networking Architectures; Pricing, Policy and Competition, Information Infrastructures for the 1990s, Harvard University, J.F. Kennedy School of Government, Nov. 1990.

McGarty, T.P., S.J. McGarty, Information Architectures and Infrastructures; Value Creation and Transfer, 19th Annual Telecommunications Policy Research Conference, Solomon's Is, MD, September, 1991.

McGarty, T.P., Communications Networks: A Morphological and Taxonomical Analysis, Columbia University, CITI Conference, October, 1991.

McGarty, T.P., Communications Network Morphological and Taxonomical Policy Implications, Telecommunications Policy Research Conference, Solomon's Island, MD, September, 1992.

McGarty, T.P., Multimedia Communications in Medicine, IEEE JSAC, November, 1992.

McGarty, T.P., S.J McGarty, Architectures et Structures de L'Information, Reseaux, No 56, pp. 119-156, December, 1992, Paris.

McGarty, T.P., Access to the Local Loop, Kennedy School of Government, Harvard University, Infrastructures in Massachusetts, March, 1993.

McGarty, T.P., Wireless Access to the Local Loop, MIT Universal Personal Communications Symposium, March, 1993.

McGarty, T.P., L.B. Ball, Integrated Network Management Systems, IEEE NOMS Conf, New Orleans, Nov. 1987.

McGarty, T.P., S.T. Treves, Multimedia Communications Applications in Health Care Services, SCAMC Conf, Washington, DC, Nov, 1990.

McGarty, T.P., Session Management in Multimedia Communications, Presented at MIT, May 2, 1990.

Pindyck, R.S., D.L Rubinfeld, Microeconomics, Macmillan (New York), 1992.

Reed, D., Putting it all Together: The Cost Structure of Personal Communications Services, OPP FCC Report Working Paper No. 28, November, 1992.

Spulber, D.F., Regulation and Markets, MIT Press (Cambridge, MA), 1990.

Tirole, J., The Theory of Industrial Organization, MIT Press (Cambridge, MA), 1988.

Weinhaus, C.L., A.G. Oettinger., Behind the Telephone Debates, Ablex (Norwood, NJ), 1988.

Winograd, T., F. Flores, Understanding Computers and Cognition, Addison Wesley (Reading, MA), 1987.