The Case for Viral Broadband

Dimitris Vyzovitis and Andrew Lippman
MIT Media Laboratory
{vyzo,lip}media.mit.edu

October 18, 2004

Abstract

A viral broadband network is a community based wireless mesh, which grows organically and utilizes the aggregate resources of the users in order to provide broadband Internet access. Effectively, the users themselves are turned into service providers, leveraging the network access of the subset of the user population. We present the case for viral broadband in the socioeconomic context of broadband Internet access.

1 Introduction

Widespread high-speed broadband Internet access is a catalyst for modern societies. However, many industry experts agree that investments and deployment of broadband infrastructure are not taking place as fast as they should [6, 12]. As a solution, the idea of bottom-up wireless broadband is actively explored [19]. Just like the development of peer-to-peer systems is changing the structure of the digital content distribution industry and represent the killer Internet application, a grass-roots, bottom-up community-based network [46] has the potential to disrupt the broadband market with a viral architecture.

A viral communications architecture[21] is one where elements are independent, scalable, and where each new element adds capacity to the system. Thus, a viral communications architecture can be adopted incrementally from a small base and gains accelerating value as the system scales. The grass-roots, bottom deployment properties of a viral communications architecture fosters innovation at the edge of the network, supporting ideas that elude the designers of centralized architectures. Such ideas, after they are technically developed in a decentralized way, can be made reliable or business-worthy, ripe for partial centralization and legitimization. The paradigm of peer-to-peer systems and applications is a prime example of this effect.

Wireless network technologies based on the 802.11 standard (WiFi) are inherently viral. They present users with cost level technology that can be used for high speed communications independent of backbone providers. Experiments like the MIT Roofnet project [42] demonstrate that the social impetus for community-based, bottom-up construction of broadband networks is supported by technology.

A viral broadband network is exactly that: a community based wireless mesh, which grows organically and utilizes the aggregate resources of the users in order to provide broadband Internet access. Effectively, the users themselves are turned into service providers, leveraging the network access of the subset of the user population.

In this paper we present the case for viral broadband, how it interacts with the economics of broadband access, and how it can stand the test of time. The evolution of broadband and the relationship with peer-to-peer systems is a prime factor in the conception of viral broadband. We outline the architecture and explore the economical aspects of viral broadband access. We show that viral broadband works both ways, for users and commercial entities, with the help of a cost distribution mechanism and support for parasitic, peer-to-peer interactions. The extrinsic requirements for the feasibility of viral broadband are modest. The base technology is already there, but successful development hinges on the open access market, reasonable behavior from service providers, and openness of the spectrum [3, 33].

The remainder of the paper is organized as follows. We review the state of the broadband market in Section 2, and derive a pricing model that applies to multi-service IP networks. We then discuss the relationship between peer-to-peer systems and the evolution of broadband in Section 3. Section 4 outlines the viral broadband architecture as an evolutionary step over wireless mesh networks. Section 5 we explore economical aspects of viral broadband using the derived pricing model. In Section 6 we reflect on the value of viral broadband as broadband decentralizing architecture, using examples from other industries, and discuss some of the technical challenges that lie ahead. The paper concludes in Section 7.

2 The State of Broadband

2.1 The Broadband Market

In the context of modern communication networks, the term broadband refers to a network in which the bandwidth can be divided and shared by multiple simultaneous signals (as for voice or data or video) [47]. As the core switching technology for various communications networks converges to packet switching, the distinction among them blurs. Without loss of generality, the *Internet Protocol* (IP) becomes the base technology for integrating voice, data, and video networks. The different types of service map to different classes of packet streams. Hence, we discuss broadband networks as IP networks, and broadband network access in terms of Internet access.

The core of the Internet has high capacity and interconnection. Internet access is provided by *Internet Service Providers* (ISP), allowing organizations and end users to connect to the network. The problem of providing high speed broadband network access for end users is known as the *last mile* problem. It pertains to the connectivity between the end user and the service provider, the first link in end user's access to the network.

At a coarse level, under the current regime of broadband access in the best effort Internet, capacity is offered at a fixed rate. We can classify end user access to the network as either symmetric or asymmetric.

In a *symmetric* end user access model, the provisioned upstream capacity C_u for a user equals the downstream capacity C_d, that is $C_u = C_d$. Symmetric access models are usually based on dedicated high capacity lines (at T1 rates or higher). An other emerging class of symmetric access is based on wireless mesh networks [42]. The idea of mesh networks is to build a wireless service provider that has a few points of contact (gateways) with the core of the network. End users are connected with each other with wireless links; packets destined for the greater Internet are routed in a multi-hop fashion through user or provider nodes, until they reach a gateway.

In an *asymmetric* end user access model, $C_u < C_d$. At present, the vast majority of end users have broadband access with an asymmetric access model [38]. The prevalent technologies used for providing asymmetric access is ADSL, which operates on top of telephone copper lines, and cable-based access, which uses the cable television network.

On top of the basic access model, broadband services are provided by content providers or value added service providers. Therefore, in the current state of affairs, we can classify the interested parties in the broadband market as:

- Base technology providers
- Infrastructure Providers
- Independent ISPs
- Content providers
- Value added service providers
- End users

Telephone companies and cable operators are infrastructure as well as service providers. They can be considered as corporate monopolies, although the legislative framework might force them to share their infrastructure with independent competitors. Independent ISPs do not own any land infrastructure. Rather, they usually utilize the infrastructure of infrastructure providers.

2.2 Economics of Broadband Access

It is interesting to examine the economics of broadband market, and how they evolve as we transition to an integrated services Internet. The point is that at the moment, the broadband Internet is limited to offering best effort data access at some fixed capacity. However the vision of broadband encompasses multi-service networks. There are multiple services simultaneously accessible to a user, including data, telephony, television, video on demand, and so on. Hence an infrastructure or service provider needs to assess its cost model, and install a charging scheme for generating revenue from the users.

In [22], the authors identify five types of cost in communication networks. These are the infrastructure cost, the connection cost for a user, the cost for expanding a network's capacity, an incremental cost for sending a packet, and a social cost. The incremental cost for transmitting a packet is very hard to quantify; when the load in the network is low, then this should be very low or zero. However, the notion of user's capacity in the core of the network is ill defined, because of statistical multiplexing and the order of core capacity compared to end user capacity. The social cost addresses the correlation between

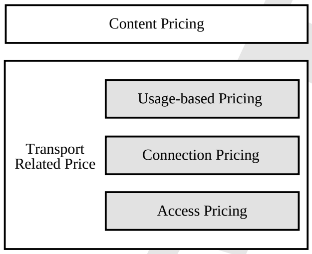

Figure 1: Components of Internet Pricing [30]

individual user's packets. A packet transmitted by a user invariably delays a packet from another user. Hence the social cost can also be considered as congestion cost [30]. It should be noted that this cost model does not fully address the application costs. For example, in a video on demand service there is a cost associated with the data transfer. Apart from the network costs there might be additional costs associated with the content, like a "rental" fee.

Traditionally, users have been charged at a fixed rate over long periods of time. While the fixed rate cost is asymptotically optimal for real-time telephony type applications [27], it is hardly the case in a multi-service network. The problem is that utility for the end user is hard to define, as there are interrelations between the various applications. In [32, 26], a four level charging scheme is suggested and analyzed. The four level scheme has components for access charge, usage charge, congestion charge, and Quality of Service (QoS) charge. Access charge is fixed, while usage charge accounts for the actual usage of the network. Congestion charge accounts for users who insist on using the network during congestion periods and should incur an additional cost. QoS charges account for users who are willing to incur a cost for guaranteeing a certain quality of service.

Charges have an associated pricing model. In [30] the authors build a layered pricing scheme, illustrated in Figure 1. The access fee is the usual fixed rate for using the network, depending on the capacity offered at the end user. Connection pricing captures resource reservation charges. Usage pricing is used to charge for services on a time, volume, or QoS basis. Finally, content pricing is an added layer, which can be determined end-to-end, to account for additional costs associated with content distribution. Depending on the service, a combination of pricing components will be present.

A multitude of pricing schemes has been proposed, the most significant ones reviewed in [11]. In general, they can be classified as static or dynamic pricing schemes. Static pricing schemes compute base costs in long time frame, while dynamic pricing schemes compute base costs within a short time-frame. An example of dynamic pricing is market-based pricing, where the individual cost for a packet is determined with an auction process amongst the users. However, as prominent researchers in the Internet community have argued [28], dynamic pricing is very hard to implement due to the complexity of accounting schemes and the difficulty of assigning utility to individual packets.

A realistic pricing scheme for multi-service Internet should compute prices locally at the edge, and facilitate receiver payments. Local computation of prices reduces the complexity of implementation and the amount of information that needs to be propagated in the network. Similarly, it provides the necessary flexibility for applying the charges to either the receiver or the sender. Receiver payments are important for content distribution applications, such as video on demand or television.

Based on these observations and the trends identified in [30], we can construct an economic model for broadband access which captures the current state of affairs and the expected evolution of broadband services in the future, with a usage-constraining policy. Specifically, consider a user accessing the network at various priorities $i = 1, \ldots, n$. For each priority, with capacity of $(C_{d,i}, C_{u,i})$, we define a fixed service rate for the priority of $c_{service,i}$, a volume cap of U_i with excess cost $c_{excess,i}$ per bit. If the network access cost is c_{access} and the content cost is $c_{content}$, then the user's charge per charging period will be

$$c = c_{access} + \sum_{i} \left(c_{service,i} + c_{excess,i} \cdot (b_i - U_i) \cdot u(b_i - U_i)\right) + c_{content} \tag{1}$$

where b_i is the number of bits transmitted or received in priority i, and $u(x)$ is the unit-step function.

For instance, in the basic broadband access model currently in place, there is a single priority for best effort service, essentially conflating $c_{service}$ and c_{access} so that $c_{service} = 0$, and $U = \infty$ as there is no volume cap. On the other hand, mobile phone operators, like Sprint, commonly offer two service classes, data and voice. Common plans provide telephony and data services with a bit cap, and charge extra for each additional bit over the limit. In this case, $U_{data}, U_{telephony}$ are both finite. Digital cable operators, for example Comcast, often provide a special service class for video on demand services

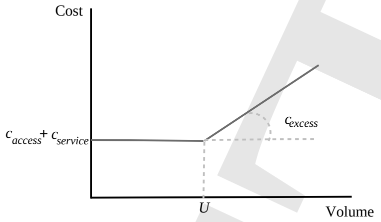

Figure 2: Pricing model for a single service, as 2-dimensional projection of the multi-service model.

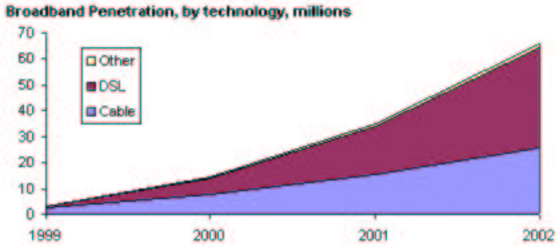

Figure 3: The early growth of Broadband [34]

(HBO), with $C_{u,video} = 0$ and $U_{video} = \infty$. Special service classes, like multicast delivery, can also easily accommodated in this model. Figure 2 illustrates the pricing model for a single service. For the case of n services, this chart is a the 2-dimensional projection.

This cost model includes the components of Figure 1, and is consistent with various charging schemes that have already been implemented for particular communication services and feasible static schemes under examination as we progress toward the integrated services Internet. We can expect usage based plans to be part of the pricing model. Similarly, we can expect multiple priority classes for various services. Finally, we can expect a fixed rate component to be present in the pricing, combined with usage based pricing or volume rate.

3 Peer-to-Peer Systems and the Evolution of Broadband

3.1 The Growth Pattern

Starting in 1999, broadband penetration has grown nearly exponentially reaching 70 million users worldwide in 2002 [34] (Figure 3). In September 2004, broadband penetration had reached 50% in the US [38]. The story of peer-to-peer file sharing networks started with Napster, which reached its heyday in 2002, after an exponential increase to some 30 million users. In contrast to conventional applications like email and the web, file sharing was the first widespread Internet application that required significant amounts of bandwidth. File sharing also represented a shift in the demands for upstream capacity, contrary to the client-server model of conventional applications. The asymmetric broadband access model is perceived as a serious roadblock in the development of peer-to-peer networks [24]. Nonetheless, the upstream capacity was sufficient for best effort reliable delivery, coupled with the number of users with sufficient upstream capacity, allowing file sharing to evolve.

Thanks to the bandwidth requirements, file sharing followed the early growth of broadband and vice versa. Following the dissolution of the original Napster, a number of decentralized systems appeared which absorbed a percentage of the user population of Napster. The wave of legal lawsuits concerning piracy and the emergence of on-line music stores have diverted significant number of users, but file sharing isn't just unauthorized music downloading. In 2004, a significant correlation between broadband access and peer-to-peer file sharing exists, as shown in [36]. According to a recent study [25] peer-to-peer applications is the single largest consumer of bandwidth on ISP networks. Over 80% of traffic in the last mile is peer-to-peer traffic. The figures are staggering when considering the traffic generated from various applications. As Figure 4 shows peer-to-peer file sharing is a killer application in the Internet.

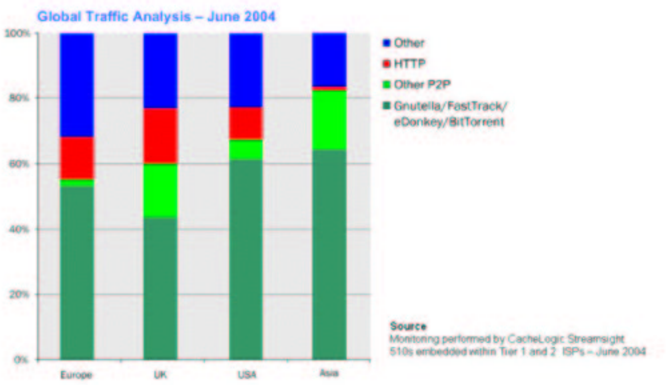

Figure 4: Statistics from global traffic analysis [25], showing the relative volume of peer-to-peer traffic

What makes peer-to-peer systems special, is the effect that they have on the perception of the Internet from end users. In the traditional client-server view of networked systems the end user is a consumer of data. In contrast, in peer-to-peer systems end users are generators of data themselves. This changes the perception of the network, from provider centric to interactive. While it is hard to predict what kind of future applications will appear in the integrated service Internet, one thing is clear: interactive peer-to-peer applications will remain a significant part of the picture and require sufficient upstream capacity to end users.

The change that peer-to-peer systems brought is typical of grass-roots disruptive applications. Despite many failures, some of these applications evolve and mature beyond the grass roots level. At this point the opportunity to centralize and legitimize these applications arises, providing new services. Maintaining the ability of the network to harbor grass roots development and innovation at the edge is essential for the broadband Internet.

There are many examples offering evidence of development of successful business ideas by centralizing disruptive applications developed at the edge of the network. For example, in the case of content distribution, major content providers like BBC are piloting peer-to-peer protocols. A study projected that by 2006 the on-line music industry will generate $2 Billion in revenue [37]. As of October 2004 the Apple iTunes service has sold over 150 million songs [35]. This industry is a direct descendant of peer-to-peer file sharing.

Another instance of a peer-to-peer application growing in popularity and morphing into a commercial enterprise is Skype. As the company is going commercial there are still open questions regarding its viability, but initial estimates are optimistic [10]. Even if the commercial offering does not turn out to be successful the effects of Skype are already felt as ISPs, such as Earthlink, are beginning to offer telephony services to their customers.

3.2 The problem of Symmetric Access

As we have noted, the prevalent broadband access model is asymmetric. The results from an unofficial survey of a small sample of top U.S. ISPs [40] are illustrated in Figure 5. The figure plots the access cost as a function of upstream capacity. Typically, upstream capacity is 4-10 times smaller than the downstream capacity. For the shake of comparison the plot also includes some “business” packages which are symmetric. Apart from technical limitations of DSL and cable-based technology, the reason for the discrepancy and high price of upstream capacity is the client-server model commonly used in assessing user needs and utility.

However, this trend is contradictory to analysis of the previous section. Peer-to-peer systems are inherently symmetric. While file-sharing, which relies on reliable best-effort delivery, is feasible an asymmetric model albeit with a performance degradation, future real-time applications will require adequate upstream capacity for end users. Examples of such applications already exist: webcams and video conferencing with quicktime generate high volume of upstream data with real time constraints.

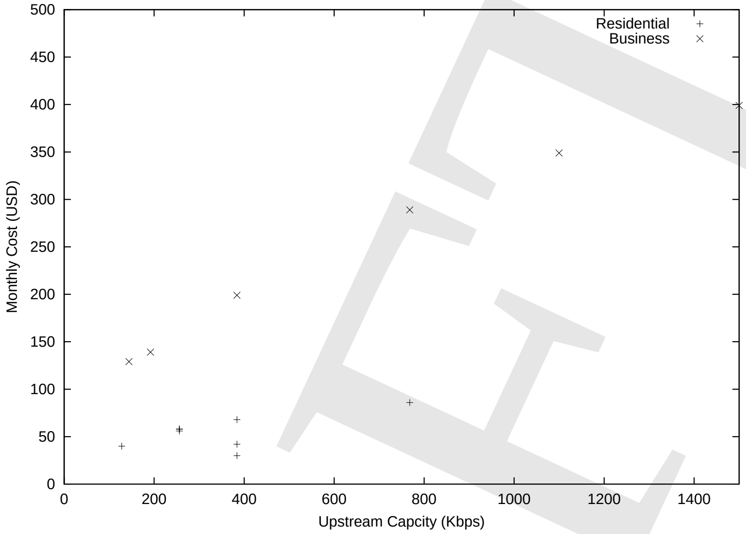

Figure 5: Monthly cost for broadband access as a function of upstream capacity from a sample of popular US providers

3.3 Peer-to-peer Systems and the Economics of Broadband

It is interesting to point out another aspect of the growth of peer-to-peer systems and broadband access. What enabled users to file share was fixed rate cost and the lack of volume caps ($U = \infty$). As a result users incurred no actual cost from file sharing.

As much as it was considered piracy, file sharing spearheaded the creation of Internet content distribution systems. The addition of $c_{content}$ in Equation 1 accounts for this effect in the cost model as end to end content pricing. If the pricing remains end-to-end, then peer-to-peer content distribution systems will continue to incur no additional cost to users, provided that a basic access service with $U = \infty$ is maintained.

However, as service providers introduce multiple priority services and volume caps the situation changes. For example, if the basic access service has a volume cap U an excess cost of $c_{excess} = a$, consider a user participating in a peer-to-peer system. If the application generates V bits of downstream data and some upstream data αV, then the user incurs an additional cost of

$$\Delta c = a \cdot ((1+\alpha)V - U) \cdot u((1+\alpha)V - U) \tag{2}$$

When the capacity bound is exceeded, Δc is nonzero, and the utility the users receives from the application changes. For the case of file sharing, this would be a disincentive for the users to participate equally, reducing α and thus exacerbating freeloading. A gnutella-like system would have significantly lower utility [1], and so would incentive based systems like BitTorrent [9].

The introduction of enhanced pricing and service models has the potential to inhibit the evolution of peer-to-peer systems. These systems are based on the resource pooling from their users. If the users incur cost by pooling their resources future peer-to-peer applications will be harder to deploy. Innovation is limited at the edge, while the big infrastructure and service providers dictate applications by controlling the costs. This situation can be avoided by maintaining a basic access best effort service without a volume cap at some reasonable capacity. Nonetheless, it is a fact that systems which require pooling of resources will have to take cost into account.

3.4 Regulation and The Evolution of the Broadband Market

Apart from pricing and cost implications, another important aspect of the evolution of the broadband market is the structure of the industry. This structure and the relationship between the players we have identified is not yet clear. However, there is a worrisome trend of promoting the interest of infrastructure providers, already established players from the telecommunications industry [6]. The broadband market is still at its incipient stages. Therefore it is important to maintain capacity for innovation at the edge of the network and thus enable the development of grassroots systems and applications.

As it is pointed out in [7], there are two schools of thought with regards to the evolution of the broadband market. At one edge of the spectrum, there is the model of allowing local phone and cable oligopolies to maintain their dominance and extend their oligopoly in the broadband market. This model relies on corporate monopolies making investments to spur broadband penetration, in effect centralizing innovation. On the other edge of the spectrum, there is the school of thought that relies on lessons learned from the traditional telecommunications industry, in terms of local access and long distance calls, and experiments in rolling broadband in other countries. In each case open access to networks and functional wholesale markets spurred lower prices and competition. The paradigm of open access is also behind the development of peer-to-peer systems.

The stated goal of the *Federal Communications Committee* (FCC) with regards to the broadband market is “to establish regulatory policies that promote competition, innovation, and investment in broadband services and facilities while monitoring progress toward the deployment of broadband services in the United States and abroad” [39]. The keyword here is *innovation*. The long history of the traditional telecommunications and other industries indicates that openness to competition is key to innovation. Unfortunately FCC seems to have aligned itself completely with the interests of established corporate players in telecommunications. The FCC has ruled that telephone companies no longer have to share their Internet networks with competitors. The result is that infrastructure providers get complete control of the market that utilizes their infrastructure. Furthermore, it makes it nearly impossible for independent service providers to establish themselves in the market, leaving end users with fewer choices and higher prices. The argument here is that the FCC should be doing exactly the opposite. As the numbers in [34] show, established monopolies attain significant market share when competing in an open access market, and therefore they do not have the basis to demand restricting it. In order to protect the interests of end users infrastructure providers should be forced into an open access market. When it was applied in the local and log distance telephony markets, following the break up of AT&T, the telecommunications industry thrived. In the contrary case [14] information access and content distribution through the Internet are in peril of being subjected to firm control.

This argument extends to End User Service Agreements as well, with regards to Internet service. A common policy for ISPs is to disallow sharing of the connection with notable exceptions like Speakeasy [43]. While end users are perceived as terminal nodes in the network graph, sharing a broadband land connection through wireless networking is a common paradigm of use. At the moment it is mostly limited to in-house sharing through wireless LANs, but this idea forms the basis of Viral Broadband. Ideally the regulatory framework would protect the right of end users to share their connection and resources. Even if it this is not the case, an open access policy would allow independent service providers to differentiate themselves and offer such provisions, adding social value.

4 Viral Broadband Architecture

As we have mentioned, a viral broadband network is a community based wireless mesh, which grows organically and utilizes the aggregate resources of the users in order to provide broadband Internet access. Effectively, the users themselves are turned into service providers, leveraging the network access of the subset of the user population. A viral broadband network accounts for multiple service providers to be involved, either wireless or wired, can aggregate multiple upstream and downstream gateways in order to provide higher available capacity to the users, and can be used to perform parasitic content distribution using locally available data. It is an evolutionary idea over wireless meshes. We first discuss the key points of a wireless mesh architecture, and then outline the architectural elements of viral broadband. We can also envision viral broadband architectures based on models different than a mesh. Nonetheless, the mesh is an understood and general enough model that suits our purposes. Finally, we discuss the case of peer-to-peer content distribution as an example system that fully utilizes the resources of a viral broadband network.

4.1 Wireless Mesh Networks

Wireless mesh networks are static ad-hoc networks for providing wireless broadband service in absence of wired infrastructure in the service area. Mobile users can roam and receive service in the covered area, but here we focus on the network backbone. There is already a number of commercial enabling technologies like the Nortel wireless mesh [41]. There are two basic approaches in the construction of a wireless mesh, as illustrated in Figure 6: ISP-centric and community-based meshes.

In an ISP-centric wireless mesh architecture (Figure 6-(a)) an ISP installs a number of router nodes or *Transit Access Points* (TAP) [17] that can reach a gateway for access to the broader Internet. Users use the transit access point as their service access point. To send a packet destined for the broader Internet, a user A forwards the packet to the TAP with which currently associated. The TAP then forwards the packet in multi-hop fashion to a gateway, in the example through TAPs

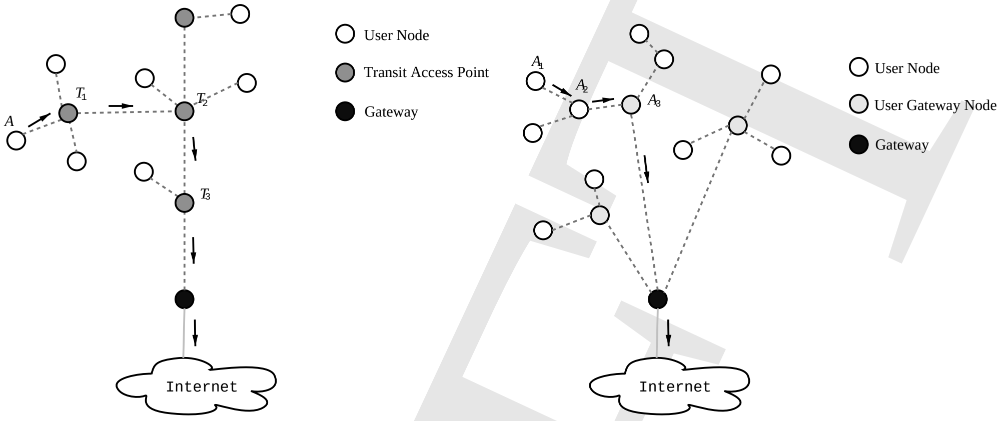

(a) (b)

Figure 6: Wireless Mesh Architecture. (a) Provider-centric model with Transit Access Points (b) Roofnet architecture with user nodes as Transit Access Points.

T_1, T_2, T_3. The ISP-centric architecture requires sufficient density for the TAPs to cover the user area. As a result, it may require significant investment from the ISP or suffer from poor coverage for some user.

Community-based wireless mesh architectures (Figure 6-(b)) address the TAP density problem by forwarding packets through user nodes until a user node with contact to a TAP is found. In Figure 6-(b), we show an example of this principle, using the architecture of Roofnet [42]. In Roofnet, a number of user nodes have line of sight to the MIT campus where gateways reside. These nodes act as TAPs for other user nodes. This is illustrated in the figure for user A_1, who goes through users A_2 and A_3 to reach a gateway.

While both approaches use multi-hop routing, there is a fundamental difference in terms of control. In an ISP-centric mesh the ISP controls the TAPs, which are essential for the operation of the entire network. In contrast, the Roofnet architecture is grass roots, since there are no ISP controlled nodes in the coverage area. A mixed architecture would include ISP-controlled TAPs inside the coverage area, but use multi-hop routing through user nodes in order to provide service for users without connectivity to a TAP.

4.2 Viral Broadband Networks

In a viral broadband network, broadband access is provided by the users themselves in a peer-to-peer fashion. The architecture is illustrated in Figure 7, including both pure and mixed models. We distinguish the following architectural elements: User nodes, User gateway nodes, TAPs, and ISP gateway nodes. TAPs and ISP gateway nodes only appear in the mixed viral model. User nodes are sources and sinks of traffic, but all generated traffic goes through the mesh. User gateway nodes have a connection to the broader Internet through an ISP (which can itself be wireless). They are both sources and sinks of traffic, but a portion of the generated traffic is not routed in the mesh. This is Internet traffic generated by the user, and can largely be routed directly through the ISP of the user.

In Figure 7-(a) we illustrate the case of capacity aggregation for a user A. A has no direct connection to the Internet. Rather, it utilizes gateways A_2 and A_5 through nodes A_1 and A_3, A_4. The two gateways contribute a capacity of a_2 and a_5 respectively. Hence the total capacity for the user is $C_A = a_2 + a_5$. The data can represent a single or multiple flows. Figure 7-(b) shows an example of user A utilizing capacity from multiple ISPs in a similar vein.

In general, consider a viral network constructed by n user nodes $A_1, \ldots, A_n$. Let the external capacity of each user be C_i, where $C_i = 0$ if user A_i does not have a direct connection to the Internet, and C_i can refer to both upstream and downstream capacity. The global external capacity of the network is then

$$C^* = \sum_i C_i \tag{3}$$

This represents the maximum capacity between the entire viral network and the broader Internet. If this capacity was equally available to all users, then each user would have an average capacity of $\frac{C}{n}$ and a peak capacity of $C_i^* = C^*$. However, in

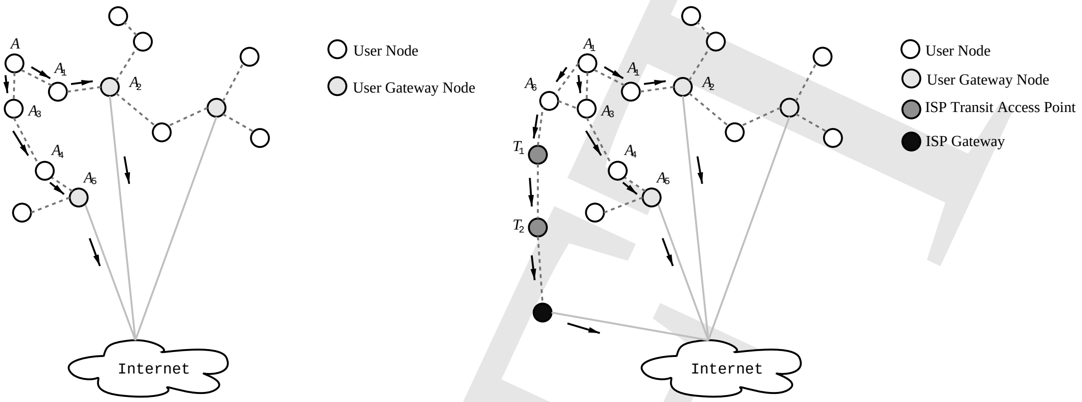

(a) (b)

Figure 7: Viral Broadband Architecture (a) Pure viral without ISP nodes (b) Mixed viral with ISP nodes

order to compute some meaningful capacity bounds, we need to take into account the internal capacity of the network. Let C_{ij} be the capacity of the link between users A_i and A_j. If the users are connected with a direct link, then C_{ij} is the capacity of that link. If there are connected through multiple hops, then C_{ij} is the minimum capacity for all hops in the optimal path from i to j, and $C_{ii} = \infty$. For simplicity, if a link is shared by more than one paths, the we divide the capacity of this link by the number of paths sharing it. Thus, we can define the per user external capacity as

$$C_i^* = \sum_j \min\{C_{ij}, C_j\} \tag{4}$$

It follows immediately that the aggregate external capacity for each user is higher than the direct external capacity:

$$\forall i, C_i^* > C_i \tag{5}$$

This implies that by adding users in the network, the peak capacity available to users increases. Of course the actual capacity available in any given moment depends on the use of the resources of the network by other users. We return to this and issues of cost in Section 5.

4.3 Viral Broadband and Peer-to-peer Content Distribution

A viral broadband network is a natural framework for performing peer-to-peer content distribution. With the content priced end-to-end, it doesn't matter where the bits come from. The user can obtain the content by an out of bands method, with regards to the content provider, and then simply authorize his use of the content – for example by obtaining a decryption key for a monetary charge from the content provider. In fact, performing the actual distribution through a peer-to-peer network reduces the cost for the content provider, as there is no transport cost incurred.

Figure 8 illustrates this scenario. User A searches the viral network for segments of the target content piece. Users B and C in the figure have parts of the content cached, enough to reconstruct the entire content. A obtains the content from B and C and contacts P, which is a provider in the broader Internet to authorize access to the content. A also becomes a cache that can support subsequent queries from other users while the content remains in the cache. Distribution efficiency can be increased by the use of multicast in the viral network. If the content is not originally available in the viral network, the user can obtain it from the broader Internet and serve as a cache for future queries in the viral network. As a result, the transport charge is eliminated for the content, increasing the efficiency of the distribution and providing gains to the users of the viral network. That is, when peer-to-peer content is distributed in its entirety in the viral network, the differential cost from peer-to-peer use is $\Delta c = 0$. In addition, content lookup and authorization also provides opportunities for value added services to be built in a viral network. A video on demand service can be easily implemented in a viral network using this approach, without requiring a massive up-front infrastructure investment to account for data transport.

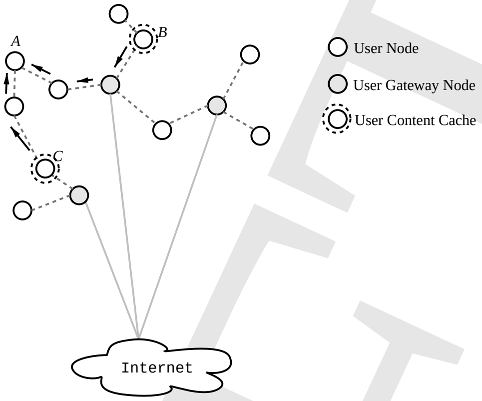

Figure 8: Peer-to-Peer Content Distribution in a Viral Network

5 Economics of Viral Broadband

5.1 Cost Distribution

From the outline of the architecture of viral broadband, it is apparent that there is a cost for providing the system with a total external capacity C^*. This cost is actually subjected to the users who have a direct connection to the Internet through an ISP. For a user A_i with a direct connection this cost is given by Equation 1, while for a user without a direct connection this cost is 0. Without any provisions for fair cost distribution amongst all users, this obviously creates a point of contention. Users with a direct connection to the Internet will be unwilling to share their resources, when doing so results in a degradation of their service while their cost remains the same.

When a basic service class with $U = \infty$ is present, then the problem of service degradation can be addressed in a simple manner. The router in the user's node can give priority to locally generated traffic, dropping packets from other users in the viral network. A TCP-friendly [13] protocol will back off, reducing the rate of traffic of flows from the viral network, and allow the user to fully utilize the connection. Similarly, a maximum rate T for flows from the viral network can be specified, ensuring that a user A_i with a direct connection always has an available capacity of $C_i - T$.

The situation is more precarious when a user is charged with volume rate in Equation 1, without a basic flat rate service class. In this case, the viral network creates additional volume of traffic b_V over the user generated traffic b, incurring an additional cost Δc:

$$\Delta c = c_{excess} \cdot ((b_V + b - U) \cdot u(b_V + b - U) - (b - U) \cdot u(b - U)) \tag{6}$$

In this case, participation in the viral network implies more than a (controlled) service degradation for a user paying for a direct Internet connection. Individual users incur cost from the community. The costs accrue similarly for different service classes (for example telephony), assuming that the viral network itself supports priority routing.

A solution to this problem is to distribute the costs according to usage. For the fixed costs, Let $A = \{A_1, \ldots, A_k\}$ be the set of users with direct connection and $A' = \{A_{k+1}, \ldots, A_n\}$ be the set of users without direct connection. In order to fairly distribute the fixed access cost of the network, each user should be charged with an access cost c^*_{access}:

$$c^*_{access} = \frac{\sum_{i \in A} c_{access,i}}{n} \tag{7}$$

Thus, each user in A' is charged with c^*_{access}, while each user $A_i \in A$ is refunded by $c_{access,i} - c^*_{access}$. The same principle applies for each service class, with the difference that users who don't use a particular class will be excluded from the costs of the class.

Usage costs can be distributed similarly. The total differential cost to be distributed among the users is $\sum_{A_j \in A} \Delta c_j$. This cost is refunded by each user $A_i \in A'$ contributing Δc^*_i:

$$\Delta c^*_i = \frac{b_i \cdot \sum_{A_j \in A} \Delta c_j}{B} \tag{8}$$

where b_i are the excess bits from user A_i, and B is the total excess volume in the system.

5.2 The Tragedy of the Commons

In addition to the fair cost distribution problem, there is another consequence of a community-based network to consider. The core of the viral broadband network is a community based mesh, where packets are routed with no additional monetary cost. However, internal routing has a social cost as it affects the routing of packets destined or sourced to the broader Internet. In conjunction with using peer-to-peer content distribution techniques, the internal traffic in the network can become quite significant and affect the ability of the users without direct connectivity to communicate with the outside world. This problem is a manifestation of the *tragedy of the commons* [16].

The tragedy of the commons is usually stated in terms of a pastry, open to all. It is to be expected that each herdsman will try to keep as many cattle as possible on the commons. A (economically) rational herdsman will seek to maximize his gain. By adding one more animal to his herd, a herdsman receives a positive utility of +1. However, doing so results on overgrazing, but the effects are shared and hence have a negative utility of $-\frac{1}{n}$, where n is the number of cattle in the pastry. The rational herdsman then concludes that adding a cattle produces positive utility, and then adds one more, and then another, and so on. However, this is the conclusion reached by each and every rational herdsman sharing the pastry. And therein lies the tragedy: Each herdsman is compelled to increase his herd without limit – in a world that is limited.

In a strict utilitarian framework, there is no technical solution for the tragedy of the commons. Nonetheless, there are some measures that can be taken to alleviate the problem in the context of viral broadband. The first measure is to establish an etiquette in the community of the users. In conjunction with the behavior of TCP-friendly protocols, users can reduce their traffic when congestion is detected. A second measure is to enforce the etiquette. Given that some accounting of bits generated by each user is necessary for cost distribution, we can define a fair share of internal bandwidth for each user. Based on this fair share, when congestion is experienced, internal routers can preferentially drop packets from users who have a history of excessive bandwidth usage.

5.3 Viral Broadband and Service Providers

The aggregation of user resources in a viral broadband network creates a point of contention with Service Providers. The system depends on the sharing of communication links to the broader Internet, provided by commercial entities to a subset of the users. As we have discussed in Section 3.4, many ISPs have End User License Agreements that prohibits sharing of the connection. Furthermore, with cost distribution schemes, it might be argued that users are reselling their capacity.

There are two ways to approach the contention between users and service providers. Firstly, we can argue that clauses prohibiting connection sharing are not really enforceable when Network Address Translation (NAT) is in use, even when ISPs can utilize machine box techniques such as the ones described in [2]. Unfortunately, this line of reasoning only creates additional tension in the relationship of users and providers, as they engage in an arms race between detecting connection sharing and disguising it.

A more reasonable approach is to allow ISPs to embrace viral broadband in enhanced service models. As we have discussed, the development of viral broadband requires a mechanism for cost distribution. While we have sketched out the solution to this problem, it is still complicated because of accounting and payment issues. The point is that accounting in a viral network presents an opportunity for providing value added services. Therefore, instead of battling viral broadband, ISPs can embrace it by providing meta-ISP services for usage accounting and payments. This is the pattern of allowing an innovative grassroots idea to evolve, and then embrace it by centralizing and legitimizing aspects of it. In a similar vein, the development of mixed viral networks provides a business opportunity. Speakeasy is already embracing the idea of end users reselling their capacity with their WiFi NetShare service [44]

Extending the argument beyond ISPs, virtually all players in the broadband market benefit from a viral network. End users receive better service at reduced cost. Value added service providers have an opportunity to offer new localized services at the user population. And base technology providers have an opportunity to develop and sell cutting edge equipment directly to end users, leveraging economies of scale.

6 Reflections on Viral Broadband

So far we have outlined the architecture of viral broadband in the context of the broadband market. We have shown that there are economic arguments for the desirability and viability of the architecture. Viral broadband builds on the social impetus of grassroots development and the technical impetus provided by wireless network technology. The community based approach allows the users to aggregate their resources and share the costs, in order to provide affordable ubiquitous broadband access. In this section we augment the argument for decentralization of broadband access with additional examples from other industries, where decentralization harbored innovation. We close with a discussion of the technical roadblocks in deployment of Viral Broadband.

6.1 Lessons from Other Industries

6.1.1 Content Distribution

From the beginning of the movie industry studios owned the rights to the content they develop. Initially, movies where distributed in a limited way through movie theaters. The development of television opened a horizon for movie and show production and distribution, as the market expanded dramatically. With a centralized model, producers generated revenue by collecting royalties for airing a movie, while television networks generated revenue from advertising.

The disruptive technology that allowed decentralization of content distribution was video recording. Video recorders offered consumers the ability to record shows and movies for playback at their leisure. It also spurred a ferocious legal battle between content producers and consumer electronic companies. The situation was resolved with the Betamax ruling. As a result video recorders where allowed to exist. This in turn spurred the creation of the video rental industry, TiVo, and netflix. The result was greater revenues for content producers and better service and prices for consumers. It is also interesting to note the evolution of the pricing model: From one-time access fee, to rental fee which allowed repeated viewing, to flat rate rental fees which has recently begun to be used by Blockbuster. The pricing models co-exist, offering the user a choice of service.

6.1.2 Long Distance Telephony

The telecommunications industry has a rich history of development. In the US, AT&T held a monopoly in the market until its break up was ordered. The market that benefited the most from was that of long distance telephony, as the companies born from the giant's breakup maintained a stranglehold of the local communications market. In contrast, the long distance telephony market is highly competitive, resulting in continuous innovation and better service and prices for the users. For a specific example of innovation that was enabled by the decentralized of the market and had an immediate impact in the service offered to end users, we can mention the development of calling cards offered by independent service providers.

6.1.3 Transportation

The development and maintenance of the highway system is essential for the operation of modern societies. However, highways have a high cost while their utility affects a subset of the population. The problem is how to subject the cost to the commuters who use a particular road artery. In the case of residential roads, the solution is easy: transfer the cost to the residents of the area through taxation. However, in application of this principle in inter-state highways or highways connecting the city to airports, this approach does not quite work as only a small subset of the population uses them. The solution to this problem was again decentralization using tolls. The pricing model evolved to variable tolls depending on the actual resource usage from commuters. This allowed some room for innovation, for example the development of Fast Lane and automatic payment systems in the US, which significantly eases charging of commuters without delaying them.

6.2 Technical Challenges

The foundation for the development of viral broadband has already been laid, both in theory and in practice. The works of [15, 20] and [29] has explored the issues of capacity bounds in wireless and ad-hoc networks, channel access schemes, and power/density relationships. In terms of practice, the base technology for viral broadband is the 802.11 standard. 802.11 enabled devices and routers are readily available and low cost. In addition, it is the technology of choice in influential projects like the MIT Roofnet and the Digital Gangetic Plains project [4], and used in commercial wireless mesh solutions. Furthermore, the development of the wimax [45] standard progresses fast, supporting additional optimism for the resolution of issues involved in terms of physical layer design.

In higher layers, there are many challenging problems to solve in order to make viral broadband possible. The way we have outlined the architecture allows us to identify some of them:

- *Routing Protocol Design*. In a basic implementation without capacity aggregation, protocols developed for ad-hoc networks [5] are already reliable and efficient enough as the use of DSR in RoofNet has shown. However, capacity aggregation, which basically amounts to multi-gateway routing, is a complicated issue.
- *Transport Protocol Design*. The basic building block of the Internet is TCP, and any newly designed transport protocol should be TCP-friendly. However, designing multi-path transport protocols is complicated. Protocols and algorithms developed for peer-to-peer content distribution, like Bullet [18] and Splitstream [8], are good starting points.
- *Multicast*. While radio communications are inherently broadcast, the design of multicast protocols for wireless networks has proved to be complicated [31]. What are suitable protocols for viral broadband?
- *Fair Sharing*. The problem of fair sharing of internal resources in the network is important. Work on fair resource sharing in peer-to-peer systems [23] provide some starting pointers in this area.
- *Accounting and Cost Distribution*. Finally, the problem of accounting and intricacies of cost distribution is critical for the viability of viral broadband systems.

7 Conclusion

In this paper we presented and supported the case for viral broadband in the socioeconomic context of broadband Internet access.

We discussed the state of the broadband market, and by identifying trends and established models in the pricing more multi-service networks, we developed a generalized pricing model, The model captures the existing conditions of Internet pricing and expected developments. The model codifies the trend toward usage-based policies.

We then discussed the evolution of broadband market, from the growth pattern to regulatory trends. We also showed the inherent relationship between peer-to-peer systems and broadband access. As a caveat for future development of broadband, we investigated the effect of usages-based pricing on peer-to-peer systems, showing the effect of provider centric pricing in disruptive applications.

In this context, we outlined the viral broadband architecture. Our analysis shows that, besides the social impetus, a viral broadband network offers concrete economic benefits for users and commercial entities alike. Viral Broadband networks are molded in a peer-to-peer fashion. We also discussed mechanisms for ensuring the survival of viral broadband in a the era of usage based pricing and integrated services. Finally, we bolstered the argument for viral broadband as a catalyst for innovation at the edge using examples from other industries.

The extrinsic requirements for the feasibility of viral broadband are modest. The base technology is already there, but successful development hinges on the open access market, reasonable behavior from service providers, and openness of the spectrum. Therefore we can conclude that viral broadband and its technical challenges is a rich area for future research that can have a resounding effect in modern societies.

References

[1] E. Adar and B. Huberman. Free riding on gnutella. *First Monday*, 5(10), October 2000. http://firstmonday.org/issues/issue5_10/adar/index.html.

[2] S. Bellovin. A technique for counting NATted hosts. In *Proc. Second Internet Measurement Workshop*, November 2002.

[3] R. Berger. Open spectrum: A path to ubiquitous connectivity. *ACM Queue*, 1(3), May 2003.

[4] P. Bhagwat, B. Raman, and D. Sanghi. Turning 802.11 inside-out. In *2nd Workshop on Hot Topics in Networks (HotNets-II)*, Cambridge, Massachusetts, 2003.

[5] S. Biswas and R. Morris. Opportunistic routing in multi-hop wireless networks. In *2nd Workshop on Hot Topics in Networks (HotNets-II)*, Cambridge, Massachusetts, 2003.

[6] Y. Braunstein. Broadband industry structure: Policy, pricing and penetration. In *Pacific Telecommunciations Conference (PTC)*, 2003.

[7] Y. Braunstein. Wrong turn in the net's future. C—Net Perspectives, July 2003. http://news.com.com/2010-1071_3-1025304.html.

[8] M. Castro, P. Druschel, A. Kermarrec, A. Nandi, A. Rowstron, and A. Singh. Splitstream: High-bandwidth multicast in a cooperative environment. In *Proc. 19th ACM Symposium on Operating Systems Principles (SOSP)*, Bolton Landing (Lake George), New York, 2003.

[9] B. Cohen. Incentives build robustness in bittorrent. In *First Workshop on the Economics of Peer-to-Peer Systems*, June 2003.

[10] D. Ewalt. Can Skype cash in on free? *Forbes Magazine*, June 2004. http://www.forbes.com/home/wireless/-2004/10/06/cx_de_1006telecom.html.

[11] M. Falkner, M. Devetsikiotis, and I. Lambadaris. An overview of pricing concepts for broadband IP networks. *IEEE Communications Review*, 3(2), 2000.

[12] C. Ferguson. The U.S. broadband problem. The Brookings Institution Policy Brief, July 2002. http://www.brookings.edu/comm/policybriefs/pb105.htm.

[13] S. Floyd and K. Fall. Promoting end-to-end congestion control in the internet. *IEEE/ACM Transactions on Networking (TON)*, 7(4), August 1999.

[14] D. Gillmor. Internet content in peril in non-competitive world. Mercury News Technology, January 2003. http://www.siliconvalley.com/mld/siliconvalley/5002309.htm.

[15] S. Gupta and P. Kumar. The capacity of wirless networks. *IEEE Transactions on Information Theory*, 46(2), March 2000.

[16] G. Hardin. The tragedy of the commons. *Science*, (162), 1968.

[17] R. Karrer, A. Sabharwal, and E. Knightly. Enabling large-scale wireless broadband: The case for TAPs. In *2nd Workshop on Hot Topics in Networks (HotNets-II)*, Cambridge, Massachusetts, 2003.

[18] D. Kostic, A. Rodriguez, J. Albrecht, and A. Vahdat. Bullet: high bandwidth data dissemination using an overlay mesh. In *Proc. 19th ACM Symposium on Operating Systems Principles (SOSP)*, Bolton Landing (Lake George), New York, 2003.

[19] W. Lehr. Wireless broadband internet. In *International Symposium on Advanced Radio Technologies*, 2003.

[20] J. Li, C. Blake, D. De Couto, H. Lee, and R. Morris. Capacity of ad hoc wireless networks. In *Proceedings of the 7th ACM International Conference on Mobile Computing and Networking (MOBICOM)*, Rome, Italy, July 2001.

[21] A. Lippman and D. Reed. Viral communications. MIT Media Laboratory Research Executive Summary, May 2003. http://dl.media.mit.edu/viral/viral.pdf.

[22] J. MacKie-Mason and H. Varian. Pricing the Internet. In B. Kahin and J. Keller, editors, *Public Access to the Internet*. MIT Press, 1995.

[23] T. Ngan, D. Wallach, and P. Druschel. Enforcing fair sharing of peer-to-peer resources. In *2nd International Workshop on Peer-to-Peer Systems (IPTPS)*, Berkeley, California, February 2003.

[24] A. Oram, editor. *Peer-to-Peer: Harnessing the power of distruptive technologies*, chapter 1. O'Reilly, 2001.

[25] A. Parker. The true picture of file sharing, 2004. http://www.cachelogic.com/research/slide1.php.

[26] C. Parris, S. Keshav, and D. Ferrari. A framework for the study of pricing in integrated networks. Technical Report TR-92-016, Berkeley, CA, 1992.

[27] I. Paschalidis and Y. Liu. Pricing in multiservice loss networks: static pricing, asymptotic optimality, and demand substitution effects. *IEEE/ACM Transactions on Networking (TON)*, 10(3), June 2002.

[28] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in computer networks: Reshaping the research agenda. *ACM Computer Communication Review*, 26, 1996.

[29] T. Shepard. *Decentralized Channel Management in Scalable Multihop Spread-Spectrum Packet Radio Networks*. PhD thesis, MIT, 1995.

[30] B. Stiller, G. Fankhauser, B. Plattner, and N. Weiler. Charging and accounting for integrated internet services - state of the art. In *Proc. INET '98*, Geneva, Switzerland, July 1998.

[31] U. Varshney. Multicast over wireless networks. *Communications of the ACM*, 45(12), December 2002.

[32] J. Warland and P. Varaiya. *High-Performance Communication Networks*. Morgan Kaufman, 1996.

[33] D. Weinberger, J. Gill, D. Hendricks, and D. Reed. Open spectrum FAQ. http://www.greaterdemocracy.org/OpenSpectrumFAQ.html.

[34] Birth of broadband. ITU Internet Reports, September 2003. http://www.itu.int/osg/spu/publications/-sales/birthofbroadband/exec_summary.html.

[35] iTunes music store downloads top 150 million songs. Apple Inc, Press release, October 2004. http://www.apple.com/pr/library/2004/oct/14itunes.html.

[36] Peer-to-peer networks in OECD countries. OECD Information Technology Outlook, October 2004. http://www.oecd.org/dataoecd/55/57/32927686.pdf.

[37] The economics of online and mobile music. Redshift Research, 2003. http://www.redshiftresearch.com/researchecon.asp.

[38] Bandwidth report, September 2004. http://www.websiteoptimization.com/bw/0409.

[39] FCC: broadband. http://www.fcc.gov/broadband.

[40] Top 23 ISPs by subscribers: Q2 2004. http://www.isp-planet.com/research/rankings/usa.html.

[41] Nortel wireless mesh solutions. http://www.nortelnetworks.com/solutions/wrlsmesh.

[42] MIT Roofnet. http://www.pdos.lcs.mit.edu/roofnet.

[43] Speakeasy. http://www.speakeasy.net.

[44] Speakeasy WiFi net share service. http://www.speakeasy.net/netshare/learnmore.

[45] WiMax forum. http://www.wimaxforum.org/home.

[46] Wireless anarchy. http://www.wirelessanarchy.com.

[47] Wordnet: A lexical database for the english language. http://wordnet.princeton.edu.