Peer-to-Peer Computing

Dejan S. Milojicic, Vana Kalogeraki, Rajan Lukose,
Kiran Nagaraja[1], Jim Pruyne, Bruno Richard,
Sami Rollins [2] ,Zhichen Xu
HP Laboratories Palo Alto
HPL-2002-57
March 8th , 2002*

E-mail: [dejan, vana, lukose, pruyne, zhichen] @ exch.hpl.hp.com, bruno_richard @ hp.com,
knagaraj @ cs. rutgers.edu, srollins @ cs.ucsb.edu

peer-to-peer,
decentralization,
self-
organization,
anonymity, cost
of ownership

The term "peer-to-peer" (P2P) refers to a class of systems and applications that employ distributed resources to perform a critical function in a decentralized manner. With the pervasive deployment of computers, P2P is increasingly receiving attention in research, product development, and investment circles. This interest ranges from enthusiasm, through hype, to disbelief in its potential. Some of the benefits of a P2P approach include: improving scalability by avoiding dependency on centralized points; eliminating the need for costly infrastructure by enabling direct communication among clients; and enabling resource aggregation.

This survey reviews the field of P2P systems and applications by summarizing the key concepts and giving an overview of the most important systems. Design and implementation issues of P2P systems are analyzed in general, and then revisited for each of the case studies described in Section 6. This survey will help people understand the potential benefits of P2P in the research community and industry. For people unfamiliar with the field it provides a general overview, as well as detailed case studies. It is also intended for users, developers, and information technologies maintaining systems, in particular comparison of P2P solutions with alternative architectures and models.

Peer-to-Peer Computing

DEJAN S. MILOJIČIĆ[1], VANA KALOGERAKI[1], RAJAN LUKOSE[1], KIRAN NAGARAJA[2], JIM PRUYNE[1], BRUNO RICHARD[1], SAMI ROLLINS[3], and ZHICHEN XU[1]

[dejan, vana, lukose, pruyne, zhichen]@exch.hpl.hp.com, bruno_richard@hp.com, knagaraj@cs.rutgers.edu, srollins@cs.ucsb.edu

[1]HP Labs, [2]Rutgers University, [3]University of California at Santa Barbara

Abstract

The term "peer-to-peer" (P2P) refers to a class of systems and applications that employ distributed resources to perform a critical function in a decentralized manner. With the pervasive deployment of computers, P2P is increasingly receiving attention in research, product development, and investment circles. This interest ranges from enthusiasm, through hype, to disbelief in its potential. Some of the benefits of a P2P approach include: improving scalability by avoiding dependency on centralized points; eliminating the need for costly infrastructure by enabling direct communication among clients; and enabling resource aggregation.

This survey reviews the field of P2P systems and applications by summarizing the key concepts and giving an overview of the most important systems. Design and implementation issues of P2P systems are analyzed in general, and then revisited for each of the case studies described in Section 6. This survey will help people understand the potential benefits of P2P in the research community and industry. For people unfamiliar with the field it provides a general overview, as well as detailed case studies. It is also intended for users, developers, and information technologies maintaining systems, in particular comparison of P2P solutions with alternative architectures and models.

1 INTRODUCTION

Peer-to-Peer (P2P) is a very controversial topic. Many experts believe that there is not much new in P2P. There is a lot of confusion: what really constitutes P2P? For example, is distributed computing really P2P or not? In addition, this field is relatively young, raising the question of whether a survey paper is needed. We believe that the confusion does warrant a thorough analysis. We have included in the paper the systems that the P2P community considers to be P2P (e.g., SETI@home), as well as those systems that have some P2P aspects (e.g., .NET my services), even if there is a divided opinion on the nature of P2P. The goals of the paper are threefold: 1) to understand what P2P is and it is not, as well as what is new, 2) to offer a thorough analysis of and examples of P2P computing, and 3) to analyze the potential of P2P computing.

The term "peer-to-peer" refers to a class of systems and applications that employ distributed resources to perform a critical function in a decentralized manner. The resources encompass computing power, data (storage and content), network bandwidth, and presence (computers, human, and other resources). The critical function can be distributed computing, data/content sharing, communication and collaboration, or platform services. Decentralization may apply to algorithms, data, and meta-data, or to all of them. This does not preclude retaining centralization in some parts of the systems and applications if it meets their requirements. Typical P2P systems reside on the edge of the Internet or in ad-hoc networks. P2P enables:

- **valuable externalities**, by aggregating resources through low-cost interoperability, the whole is made greater than the sum of its parts

---

- **lower cost of ownership and cost sharing**, by using existing infrastructure and by eliminating and distributing the maintenance costs

- **anonymity/privacy**, by incorporating these requirements in the design and algorithms of P2P systems and applications, and by allowing peers a greater degree of autonomous control over their data and resources

However, P2P also raises some security concerns for users and accountability concerns for IT. In general it is still a technology in development where it is hard to distinguish useful from hype and new from old. In the rest of the paper we evaluate these observations in general as well as for specific P2P systems and applications.

P2P gained visibility with Napster's support for music sharing on the Web [Napster 2001] and its law suit with the music companies over digital rights management. However, it is increasingly becoming an important technique in various areas, such as distributed and collaborative computing both on the Web and in ad-hoc networks. P2P has received the attention of both industry and academia. Some big industrial efforts include the P2P Working Group, led by many industrial partners such as Intel, HP, Sony, and a number of startup companies; and JXTA, an open-source effort led by Sun. There are a number academic events dedicated to P2P, such as the International Workshop on P2P Computing, Global and P2P Computing on Large Scale Distributed Systems, International Conference on P2P Computing, and O'Reilly P2P and Web Services Conference. There are already a number of books published [Oram 2000, Barkai 2001, Miller, 2001, Moore and Hebeler 2001, Fatah and Fatah 2002], and a number of theses and projects in progress at universities, such as Chord [Stoica et al 2001], OceanStore [Kubiatowitz et al. 2000], PAST [Druschel and Rowstron 2001], CAN [Ratnasamy 2001], and FreeNet [Clark 1999].

There are several of the definitions of P2P that are being used by the P2P community. The Intel P2P working group defines it as "the sharing of computer resources and services by direct exchange between systems" [p2pwg, 2001]. Alex Weytsel of Aberdeen defines P2P as "the use of devices on the internet periphery in a non-client capacity" [Veytsel, 2001]. Ross Lee Graham defines P2P through three key requirements: a) they have an operational computer of server quality; b) they have an addressing system independent of DNS; and c) they are able to cope with variable connectivity [Graham 2001]. Clay Shirky of O'Reilly and Associate uses the following definition: "P2P is a class of applications that takes advantage of resources – storage, cycles, content, human presence – available at the edges of the Internet. Because accessing these decentralized resources means operating in an environment of unstable connectivity and unpredictable IP addresses, P2P nodes must operate outside the DNS system and have significant or total autonomy from central servers" [Shirky 2001]. Finally, Kindberg defines P2P systems as those with independent lifetimes [Kindberg 2002].
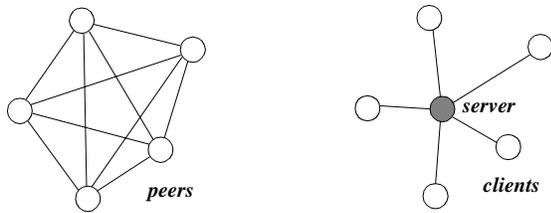
Figure 1: Simplified, High-Level View of Peer-to-Peer versus Centralized (Client-Server) Approach.

In our view, P2P is about sharing: giving to and obtaining from the peer community. A peer gives some resources and obtains other resources in return. In the case of Napster, it was about offering music to the rest of the community and getting other music in return. It could be donating resources for a good cause, such as searching for extraterrestrial life or combating cancer, where the benefit is obtaining the satisfaction of helping others. P2P is also a way of implementing systems based on the notion of increasing the decentralization of systems, applications, or simply algorithms. It is based on the principles that the world will be connected and widely distributed and that it will not be possible or desirable to leverage everything off of centralized, administratively managed infrastructures. P2P is a way to leverage vast amounts of computing power, storage, and connectivity from personal computers distributed around the world.

Assuming that "peer" is defined as "like each other," a P2P system then is one in which autonomous peers depend on other autonomous peers. Peers are autonomous when they are not wholly controlled by each other or by the same authority, e.g., the same user. Peers depend on each other for getting information, computing resources, forwarding requests, etc. which are essential for the functioning of the system as a whole and for the benefit of all peers. As a result of the autonomy of peers, they cannot necessarily trust each other and rely completely on the behavior of other peers, so issues of scale and redundancy become much more important than in traditional centralized or distributed systems.

Conceptually, P2P computing is an alternative to the centralized and client-server models of computing, where there is typically a single or small cluster of servers and many clients (see Figure 1). In its purest form, the P2P model has no concept of server; rather all participants are peers. This concept is not necessarily new. Many earlier distributed systems followed a similar model, such as UUCP [Nowitz 1978] and switched networks [Tanenbaum 1981]. The term P2P is also not new. In one of its simplest forms, it refers to the communication among the peers. For example, in telephony users talk to each other directly once the connection is estab-
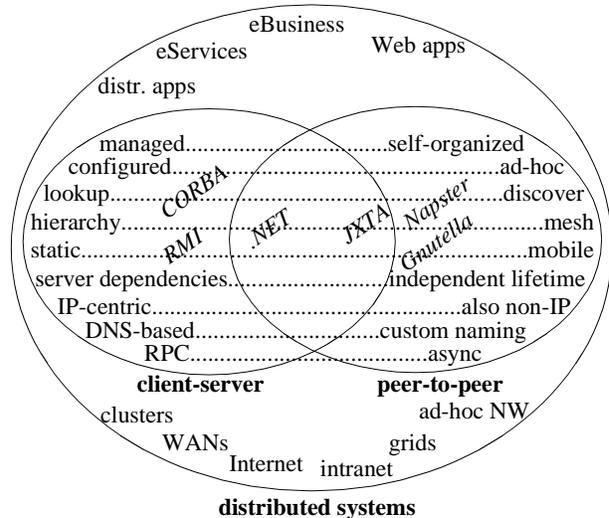

Figure 2: Peer-to-Peer versus Client-Server. *There is no clear border between a client-server and a P2P model. Both models can be built on a spectrum of levels of characteristics (e.g., manageability, configurability), functionality (e.g., lookup versus discovery), organizations (e.g., hierarchy versus mesh), components (e.g., DNS), and protocols (e.g., IP), etc. Furthermore, one model can be built on top of the other or parts of the components can be realized in one or the other model. Finally, both models can execute on different types of platforms (Internet, intranet, etc.) and both can serve as an underlying base for traditional and new applications. Therefore, it should not be a surprise that there is so much confusion about what P2P is and what it is not. It is extremely intertwined with existing technologies [Morgan 2002].*

lished, in a computer networks the computers communicate P2P, and in games, such as Doom, players also interact directly. However, a comparison between client-server and P2P computing is significantly more complex and intertwined along many dimensions. Figure 2 is an attempt to compare some aspects of these two models.

The P2P model is quite broad and it could be evaluated from different perspectives. Figure 3 categorizes the scope of P2P development and deployment. In terms of development, platforms such as JXTA provide an infrastructure to support P2P applications. Additionally, developers are beginning to explore the benefit of implementing various horizontal technologies such as distributed computing, collaborative, and content sharing software using the P2P model rather than more traditional models such as client-server. Applications such as file sharing and messaging software are being deployed in a number of different vertical markets. Section 2.3 provides a more thorough evaluation of P2P markets and Section 5 describes the horizontal technologies in more detail.

Perspective	Comparison								
	What is New		**What is Not New**		**Well-Known Examples**	**Enabler**	**Enabling**	**Alternatives**	**Paper Roadmap**
Historical/Evolutionary (Computing)	computing on the edge of the Internet	scalability, availability, security, connectivity	distributed scheduling	concept, applications, distribution	SETI@home	ubiquitous computing/communication	use vast available computer power at home and office	classic NW, clusters, Grids	Sections 5.2, 6.1, 6.2
Cultural/Sociological (Content sharing/Services)	direct sharing (privacy, anonymity)		decentralization		Napster	broad Internet connectivity	user-to-user exchange, minimal broker engagement	client-server, B2B Web Services	Sections 5.3, 6.5, 6.6
Communication/ Collaboration	dealing with disconnection		ad-hoc NW, disconnected operation		AOL Chat, Groove parasitic NW	new NWs: wireless, broadband	improved communication and collaboration	Lotus Notes, NetMeeting	Sections 5.4, 6.3, 6.4
Architectural	cost of ownership		P2P concept and applications		JXTA, .NET	increased component decentralization	larger scale, better accessibility	CORBA, RMI, other middleware	Sections 5.5, 6.7, 6.8
Algorithms/ Programming Model	particular algorithms		distributed state algor. in general		Gnutella	decentralized state	improved scalability, availability, and anonymity	client-server	Section 5.3

Table 1. Various P2P Perspectives Illustrating What is and What is Not New in P2P.

The following three lists, which are summarized in Table 1. are an attempt to define the nature of P2P, what is new in P2P, and what is not new in P2P.

P2P is concerned with:

- The historical evolution of computing in general and the Internet in particular; computing at the edge of the Internet. (e.g., SETI@home and other distributed computing systems)
- Sociological aspects of sharing of content (e.g., Napster and other file/content sharing systems)
- Technological improvements to networks and communication advances (e.g., wireless networks, handheld devices enabling better collaboration and communication)
- P2P software architectures (e.g., JXTA or .NET)
- Deployed P2P algorithms (e.g., Gnutella, FreeNet).

New aspects of P2P include:

- Technology requirements
  - scale of deployed computers
  - ad-hoc connectivity
  - security (e.g., crossing the firewall)

- Architectural requirements
  - availability in particular of the scale of the future number of computers
  - scalability of future systems world-wide as well as of embedded systems
  - privacy and anonymity (beyond the Web)
- Economy requirements
  - cost of ownership
  - pervasive use (home versus business – expectations about quality and guarantees versus best effort)

Aspects of P2P that are not new include:

- Concept and applications (e.g., telephony, networks, servers)
- Decentralization for scalability and availability (e.g., distributed systems in general, replication)
- Distributed state management (e.g., work of Barak and many others)
- Disconnected operations in general (e.g., Coda and many others)
- Distributed scheduling algorithms (e.g., on clusters and grids)
- Scalability (WWW and Web services in particular)
- Ad-hoc networks
- eBusiness
- Algorithms (many P2P and distributed algorithms already exist)

Paper Organization and Intended Audience

The paper is organized as follows. Section 2 provides background on P2P. Section 4 describes characteristics of P2P solutions. Section 3 presents P2P components and algorithms. Section 5 classifies P2P systems into five categories and describes a few representatives for each category. In Section 6, we present more detailed case studies. Section 7 presents lessons learned from analyz-

Figure 3: P2P Solutions. *P2P can be classified into interoperable P2P platforms, applications of the P2P technology, and vertical P2P applications.*

ing the P2P area. Finally, in the last section, we summarize the paper and describe opportunities for further research. In Appendix A, we compare P2P systems with alternative solutions.

This paper is intended for people new to the field of P2P as well as for experts. It is intended for users, developers, and IT personnel. This first section provided a brief overview of the field. Readers interested in learning about the field in general should read Sections 2 through 5. Experts can benefit from the detailed case studies in Section 6. Perspectives of users, developers, and IT personnel are addressed in Section 7.3. We assume that the readers have a general knowledge of computing systems.

2 OVERVIEW

P2P is frequently confused with other terms, such as traditional distributed computing [Coulouris et al. 2001], grid computing [Foster and Kesselman 1999], and ad-hoc networking [Perkins 2001]. To better define P2P, this section introduces P2P goals, terminology, and taxonomies.

2.1 Goals

As with any computing system, the goal of P2P systems is to support applications that satisfy the needs of users. Selecting a P2P approach is often driven by one or more of the following goals.

- **Cost sharing/reduction**. Centralized systems that serve many clients typically bear the majority of the cost of the system. When that main cost becomes too large, a P2P architecture can help spread the cost over all the peers. For example, in the file-sharing space, the Napster system enabled the cost sharing of file storage, and was able to maintain the index required for sharing. In the end, the legal costs of maintaining the index became too large, and so more radical P2P systems such as Gnutella were able to share the costs even further. Much of the cost sharing is realized by the utilization and aggregation of otherwise unused resources (e.g. SETI@home), which results both in net marginal cost reductions and a lower cost for the *most* costly system component. Because peers tend to be autonomous, it is important for costs to be shared reasonably equitably.
- **Improved scalability/reliability**. With the lack of strong central authority for autonomous peers, improving system scalability and reliability is an important goal. As a result, algorithmic innovation in the area of resource discovery and search has been a clear area of research, resulting in new algorithms for existing systems [REFS], and the development of new P2P plat-

forms (such as CAN [Ratnasamy 2001], Chord [Stoica 2001], and PAST [Rowstron 2001]). These developments will be discussed in more detail in Section 3.

- **Resource aggregation and interoperability**. A decentralized approach lends itself naturally to aggregation of resources. Each node in the P2P system brings with it certain resources such as compute power or storage space. Applications that benefit from huge amounts of these resources, such as compute-intensive simulations or distributed file systems, naturally lean toward a P2P structure to aggregate these resources to solve the larger problem. Distributed computing systems, such as SETI@Home, distributed.net, and Endeavours are obvious examples of this approach. By aggregating compute resources at thousands of nodes, they are able to perform computationally intensive functions. File sharing systems, such as Napster, Gnutella, and so forth, also aggregate resources. In these cases, it is both disk space to store the community's collection of data and bandwidth to move the data that is aggregated. Interoperability is also an important requirement for the aggregation of diverse resources.
- **Increased autonomy**. In many cases, users of a distributed system are unwilling to rely on any centralized service provider. Instead, they prefer that all data and work on their behalf be performed locally. P2P systems support this level of autonomy simply because they require that the local node do work on behalf of its user. The principle example of this is the various file sharing systems such as Napster, Gnutella, and FreeNet. In each case, users are able to get files that would not be available at any central server because of licensing restrictions. However, individuals autonomously running their own servers have been able to share the files because they are more difficult to find than a server operator would be.
- **Anonymity/privacy**. Related to autonomy is the notion of anonymity and privacy. A user may not want anyone or any service provider to know about his or her involvement in the system. With a central server, it is difficult to ensure anonymity because the server will typically be able to identify the client, at least by Internet address. By employing a P2P structure in which activities are performed locally, users can avoid having to provide any information about themselves to anyone else. FreeNet is a prime example of how anonymity can be built into a P2P application. It uses a forwarding scheme for messages to ensure that the original requestor of a service cannot be tracked. It increases anonymity by using probabilistic algorithms so that origins cannot be easily tracked by analyzing network traffic.

5

- **Dynamism**. P2P systems assume that the computing environment is highly dynamic. That is, resources, such as compute nodes, will be entering and leaving the system continuously. When an application is intended to support a highly dynamic environment, the P2P approach is a natural fit. In communication applications, such as Instant Messaging, so-called "buddy lists" are used to inform users when persons with whom they wish to communicate become available. Without this support, users would be required to "poll" for chat partners by sending periodic messages to them. Likewise, distributed computing applications such as distributed.net and SETI@Home must adapt to changing participants. They, therefore, must re-issue computation jobs to other participants to ensure that work is not lost if earlier participants drop out of the network while they were performing a computation step.

- **Enabling ad-hoc communication and collaboration**. Related to dynamism is the notion of supporting ad-hoc environments. By ad hoc, we mean environments where members come and go based perhaps on their current physical location or their current interests. Again, P2P fits these applications because it naturally takes into account changes in the group of participants. P2P systems typically do not rely on established infrastructure — they build their own, e.g., logical overlay in CAN and PAST.

2.2 Terminology

The following terminology is used in the taxonomies in Section 2.3 and in the comparisons between P2P and its alternatives in Section 7.1 and Appendix A.

- *Centralized systems* represent single-unit solutions, including single- and multi-processor machines, as well as high-end machines, such as supercomputers and mainframes.

- *Distributed systems* are those in which components located at networked computers communicate and coordinate their actions only by passing messages [Couloris, et al. 2001].

- *Client* is informally defined as an entity (node, program, module, etc.) that initiates requests but is not able to serve requests. If the client also serves the request, then it plays the role of a server.

- *Server* is informally defined as an entity that serves requests from other entities, but does not initiate requests. If the server does initiate requests, then it plays the role of a client. Typically, there are one or a few servers versus many clients.

- *Client-Server model* represents the execution of entities with the roles of clients and servers. Any entity in a system can play both roles but for a different purpose, i.e. server and client functionality residing on separate nodes. Similarly an entity can be a server for one kind of request and client for others.

- *Peer* is informally defined as an entity with capabilities similar to other entities in the system.

- *P2P model* enables peers to share their resources (information, processing, presence, etc.) with at most a limited interaction with a centralized server. The peers may have to handle a limited connectivity (wireless, unreliable modem links, etc.), support possibly independent naming, and be able to share the role of the server [Oram, 2000]. It is equivalent to having all entities being client and servers for the same purpose.

Other terms frequently associated with P2P include:

- *Distributed computing, which* is defined as "a computer system in which several interconnected computers share the computing tasks assigned to the system" [IEEE 1990]. Such systems include *computing clusters*, *Grids* (see below), and *global computing systems* gathering computing resources from individual PCs over the Internet. We will use the term *Distributed Computing* to refer to systems that have inherent P2P properties.

- *Grid computing*, which is defined as "coordinated resource sharing and problem solving in large, multi-institutional virtual organization." [Foster and Kesselman 1999]. More specifically, a Grid is an infrastructure for globally sharing compute-intensive resources such as supercomputers or computational clusters. As far as transparency is concerned, Grid computing is orthogonal to P2P distributed computing systems.

- *Ad-hoc communication, which is defined as a system that* enables communication to take place without any preexisting infrastructure in place, except for the communicating computers. These computers form an ad-hoc network. This network and associated computers take care of communication, naming, and security. P2P systems can be used on top of an ad-hoc communication infrastructure.

There are many examples of distributed systems, at various scales, such as the Internet, wide-area networks, intranets, local-area networks, etc. Distributed system components can be organized in a P2P model or in a client-server model. (We believe that other models, such as three-tier and publish-subscribe, can be mapped onto client-server). Typical client examples include Web browsers (e.g., Netscape Communicator or Internet Explorer),
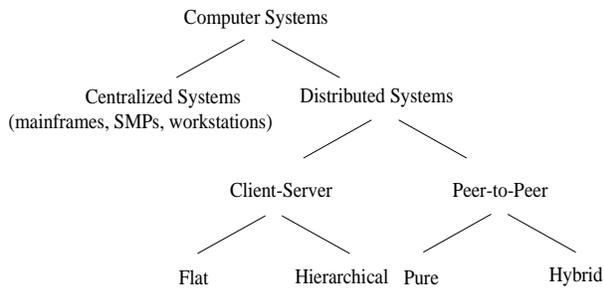
Figure 4: A Taxonomy of Computer Systems.

file system clients, DNS clients, CORBA clients, etc. Server examples include name servers (DNS [Albitz and Liu 2001, Mockapetris 1989], LDAP [Howes and Smith 2001], etc.]), distributed file servers (NFS [Sandberg et al. 1985], AFS [Howard et al. 1988], CORBA Object Request Brokers [OMG 1996], HTTP Server, authentication server, etc. Client-server model examples include CORBA, RMI [Wolrath et al 1996], and other middleware [Bernstein 1996; Britton 2000]). Peer examples include computers in a network that serve a similar role. There are numerous examples of the P2P model throughout the paper.

2.3 P2P Taxonomies

A taxonomy of computer systems from the P2P perspective is presented in Figure 4. All computer systems can be classified into centralized and distributed. Distributed systems can be further classified into the client-server model and the P2P model. The client-server model can be flat where all clients only communicate with a single server (possibly replicated for improved reliability), or it can be hierarchical for improved scalability. In a hierarchal model, the servers of one level are acting as clients to higher level servers. Examples of a flat model include traditional middleware solutions, such as object request brokers and distributed objects. Examples of a hierarchical model include DNS server and mounted file systems.

The P2P model can either be *pure* or it can be *hybrid*. In a pure model, there does not exist a centralized server.
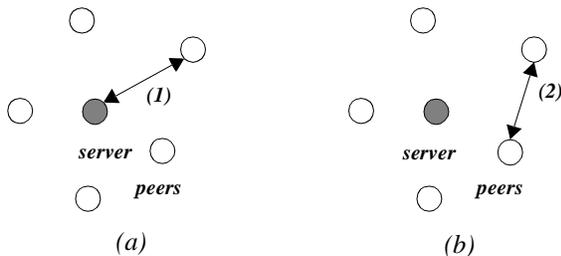
Figure 5: Hybrid Peer-to-Peer Model. *(1) Initial communication with a server, e.g., to obtain the location/identity of a peer, followed by (2) direct communication with a peer.*
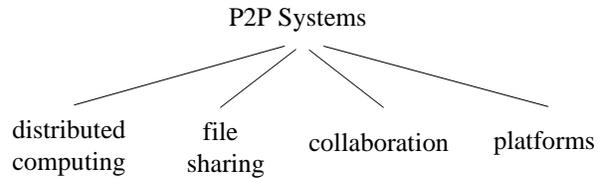
Figure 6: A Taxonomy of P2P Systems.

Examples of a pure P2P model include Gnutella and Freenet. In a hybrid model, a server is approached first to obtain meta-information, such as the identity of the peer on which some information is stored, or to verify security credentials (see Figure 5 (a)). From then on, the P2P communication is performed (see Figure 5 (b)). Examples of a hybrid model include Napster, Groove, Aimster, Magi, Softwax, and iMesh. There are also intermediate solutions where with SuperPeers, such as KaZaa. SuperPeers contain some of the information that others may not have. Other peers typically lookup information at SuperPeers if they cannot find it otherwise.

Figure 6 presents a coarse taxonomy of P2P systems. We classify P2P systems into d*istributed computing* (e.g., SETI@home, Avaki, Entropia), f*ile sharing* (e.g., Napster, Gnutella, Freenet, Publius, Free Haven), *collaboration* (e.g., Magi, Groove, Jabber), and *platforms* (e.g., JXTA and .NET My Services). Section 6 contains a description of eight case studies of P2P systems according to the taxonomy presented in Figure 6.

In Figure 7, we present a classification of various P2P systems according to the taxonomy presented in Figure 6. This figure demonstrates that certain P2P systems emphasize different aspects along the taxonomy dimensions (computing, storage, communication), whereas the platforms support all of these dimensions.
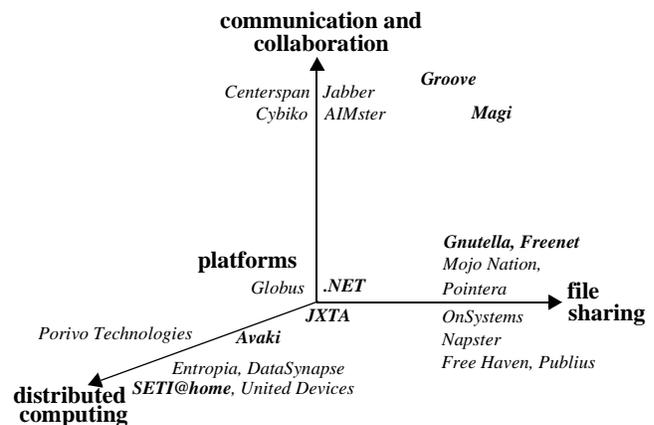
Figure 7: A Classification of P2P Systems Based on the Taxonomy in Figure 6. *Systems in bold-italic are case studies described in detail in Section 6.*
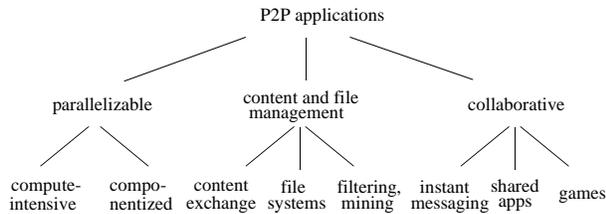
Figure 8: A Taxonomy of P2P Applications.

Application Taxonomy. Three main classes of P2P applications have emerged: parallelizable, content and file management, and collaborative (see Figure 8).

- **Parallelizable.** Parallelizable P2P applications split a large task into smaller sub-pieces that can execute in parallel over a number of independent peer nodes. Most implementations of this model have focused on compute-intensive applications. The general idea behind these applications is that idle cycles from any computer connected to the Internet can be leverage to solve difficult problems that require extreme amounts of computation. Most often, the same task is performed on each peer using different sets of parameters. Examples of implementations include searching for extraterrestrial life [SETI@Home, 2001], code breaking, portfolio pricing, risk hedge calculation, market and credit evaluation, and demographic analysis. Section 5.2 presents systems supporting this class of application. Componentized applications have not yet been widely recognized as P2P. However, we envision applications that can be built out of finer-grain components that execute over many nodes in parallel. In contrast to compute-intensive applications that run the same task on many peers, componentized applications run different components on each peer. Examples include Workflow, JavaBeans, or Web Services in general.

- **Content and file management.** Content and file management P2P applications focus on storing information on and retrieving information from various peers in the network. The model that popularized this class of application is the content exchange model. Applications like Napster [Napster 2001] and Gnutella [Gnutella 2001] allow peers to search for and download files, primarily music files, that other peers have made available. For the most part, current implementations have not focused much on providing reliability and rely on the user to make intelligent choices about the location from which to fetch files and to retry when downloads fail. They focus on using otherwise unused storage space as a server for users. These applications must ensure reliability by using more traditional database techniques such as replication. A number of research projects have explored the foundations of P2P file systems [Ratnasamy et al 2001, Bolosky et al 2000, Kubiatowicz et al 2000, Rowstron and Druschel 2001, Gribble et al 2001, Stoica et al 2001]. Finally, filtering and mining applications such as OpenCOLA [OpenCOLA 2001] and JXTA Search [Waterhouse et al. 2002] are beginning to emerge. Instead of focusing on sharing information, these applications focus on collaborative filtering techniques that build searchable indices over a peer network. A technology such as JXTA Search can be used in conjunction with an application like Gnutella to allow more up-to-date searches over a large, distributed body of information.

- **Collaborative.** Collaborative P2P applications allow users to collaborate, in real time, without relying on a central server to collect and relay information. Instant messaging is one subclass of this class of application. Services such as Yahoo!, AOL, and Jabber instant messaging have become popular among a wide variety of computer users [Strom 2001]. Similarly, shared applications that allow people (e.g., business colleagues) to interact while viewing and editing the same information simultaneously, yet possibly thousands of miles apart, are also emerging. Examples include Buzzpad [www.buzzpad.com] and distributed Power-Point [Rice and Mahon 2000]. Games are a final type of collaborative P2P application. P2P games are hosted on all peer computers and updates are distributed to all peers without requiring a central server. Example games include NetZ 1.0 by Quazal [www.quazal.com], Scour Exchange by CenterSpan, Descent [www.planetdescent.com], and Cybiko.

P2P Target Environments. The target environments for P2P consist of the *Internet, intranets,* and *ad-hoc networks*. P2P systems connected to Internet support connections in the spectrum from dialup lines to broadband (DSL). The underlying architecture can rely on *personal home computers, corporate desktops*, or *personal mobile computers* (laptops and handhelds).

The most frequent environment is personal home computers connected to the Internet. The early P2P systems in this environment were primarily used for content sharing. Examples include Napster, Gnutella, and Aimster. Distributed computing in a P2P fashion also started on desktop computers on the Internet such as SETI@home, but it also gained acceptance in Intranets, such as in DataSynapse [Intel 2001, DataSynapse 2001].

Ad-hoc networks of handhelds are only recently becoming available (e.g., Endeavors Technologies Magi) dedicated for collaborative computing, but similar platforms are expected to follow with a wider deployment of handheld computers and wireless networks.
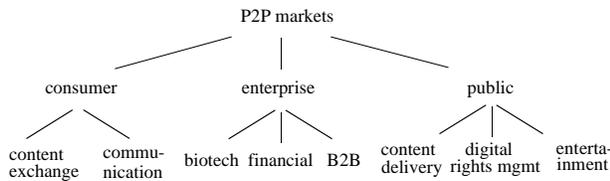
Figure 9: A Taxonomy of P2P Markets. *There are also other markets that can map to some of the presented ones, such as military, education, scientific, etc.*

Future environments may include variations of existing deployment scenarios, such as using corporate desktops for content sharing, e.g., for distributed Internet data centers, or handheld computers for resource aggregation (content sharing or distributed computing). Furthermore, Internet2 (http://www.internet2.edu/) may offer more enabling support for P2P systems and applications.

P2P Markets. There are three main markets for P2P: *consumer*, *enterprise*, and *public*. *Consumer space* encompasses the use of P2P for personal use, such as music and content sharing, instant messaging, email, and games. Application examples include music sharing and communication. Examples of companies in this space include Napster and Gnutella. *Enterprise space* P2P applications include biotech, financial, traditional IT solutions, and B2B. Example companies include Data Synapse, Information Architects, and WorldStreet. Finally, the public class of applications includes the following types of applications: information sharing, digital rights management, and entertainment. Centerspan, AIM, and Scour deliver music and video on broadband using P2P technology.

According to Philips, markets can be classified based on user activities into those performed @work, @play, and @rest [Nethisinghe 200]. We extended this perspective to map user activities to markets (consumer, enterprise, and public) (see Table 2). The insight is that P2P permeates a number of user activities and markets.

type of activity	scope		
	consumer	**enterprise**	**public**
@work	collaboration, communication	distributed computing, storage, communication, collaboration	communication, digital rights mgmt
@play	games	HR-sponsored events	digital media digital experience
@rest	music sharing	content consumption	instant messaging

Table 2. P2P Markets versus P2P Applications.
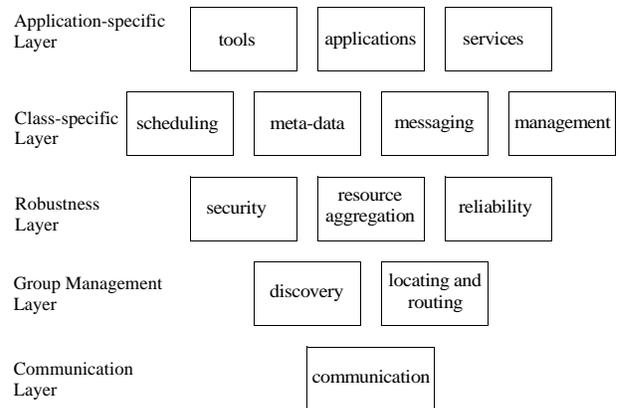
Figure 10: An Informal P2P System Architecture. *The order of components does not strictly follow layering.*

3 COMPONENTS AND ALGORITHMS

This section introduces components of P2P systems and algorithms used in P2P.

3.1 Infrastructure Components

Figure 10 illustrates an abstract P2P architecture. In this section, we discuss the functions of each component and look at some of the tradeoffs involved in making component implementation decisions.

Communication. The P2P model covers a wide spectrum of communication paradigms. At one end of the spectrum are desktop machines mostly connected via stable, high-speed links over the Internet [Saroiu et al. 2002]. At the other end of the spectrum, are small wireless devices such as PDAs or even sensor-based devices that are connected in an ad-hoc manner via a wireless medium. The fundamental challenge of communication in a P2P community is overcoming the problems associated with the dynamic nature of peers. Either intentionally (e.g., because a user turns off her computer) or unintentionally (e.g., due to a, possibly dial-up, network link failing) peer groups frequently change. Maintaining application-level connectivity in such an environment is one of the biggest challenges facing P2P developers.

Group Management. Peer group management includes discovery of other peers in the community and location and routing between those peers. Discovery of peers can be highly centralized such as in Napster [Napster 2001], highly distributed such as in Gnutella [Gnutella 2001], or somewhere in between. A number of factors influence the design of discovery algorithms. For example, mobile, wireless devices can discover other peers based upon

9

their range of communication [Roman et al. 2001]. Protocols built for desktop machines often use other approaches such as centralized directories.

Location and routing algorithms generally try to optimize the path of a message traveling from one peer to another. While deployed systems such as Napster and Gnutella try to optimize factors such as underlying network latency, the research in this space takes a more systematic approach as discussed in Section 3.2.

Robustness. There are three main components that are essential to maintaining robust P2P systems: security, resource aggregation, and reliability. Security is one of the biggest challenges for P2P infrastructures. A benefit of P2P is that it allows nodes to function as both a client and a server. However, transforming a standard client device into a server poses a number of risks to the system. Only trusted or authenticated sources should have access to information and services provided by a given node. Unfortunately, the security requirement either requires potentially cumbersome intervention from the user, or interaction with a trusted third party. Centralizing the task of security is often the only solution even though it voids the P2P benefit of a distributed infrastructure. For a more detailed discussion of P2P security concerns see Section 4.8.

The P2P model provides the basis for interacting peers to aggregate resources available on their systems. Classifying the architecture of a P2P resource aggregation component is difficult because there are a wide variety of resources that may be aggregated across peers. On the one hand, resources include files or other content residing on a computer. A wide variety of file sharing systems have addressed the problem of aggregating this type of resource. On the other hand, resources can be defined in terms of the resources available on a given peer device such as CPU processing power, bandwidth, energy, and disk space.

Reliability in P2P systems is a hard problem. The inherently distributed nature of peer networks makes it difficult to guarantee reliable behavior. The most common solution to reliability across P2P systems is to take advantage of redundancy. For example, in case of compute-intensive applications upon a detection of a failure the task can be restarted on other available machines. Alternatively, the same task can be initially assigned to multiple peers. In file sharing applications, data can be replicated across many peers. Finally, in messaging applications, lost messages can be resent or can be sent along multiple paths simultaneously.
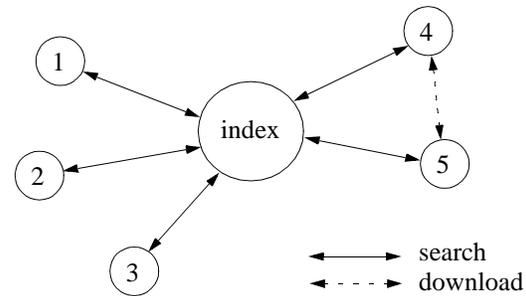
Figure 11: Central Index Algorithm.

Class-Specific. While the components discussed so far are applicable to any P2P architecture, application-specific components abstract functionality from each class of P2P application. Scheduling applies to parallelizable or compute-intensive applications. Compute-intensive tasks are broken into pieces that must be scheduled across the peer community. Metadata applies to content and file management applications. Metadata describes the content stored across nodes of the peer community and may be consulted to determine the location of desired information. Messaging applies to collaborative applications. Messages sent between peers enable communication. Management supports managing the underlying P2P infrastructure.

Application-Specific. Tools, applications, and services implement application-specific functionality, which correspond to certain P2P applications running on top of the underlying P2P infrastructure. It corresponds to specific cases of distributed scheduling (e.g.scientific, financial, biotechnology, etc.), content and file sharing (e.g., music MP3 file exchange), or specific applications running on top of collaborative and communication systems, such as calendaring, notes, messaging, and chatting.

3.2 Algorithms

This section overviews three common P2P algorithms and then compares their implementations in a few P2P systems.

Centralized directory model. This model was made popular by Napster. The peers of the community connect to a central directory where they publish information about the content they offer for sharing (see Figure 11). Upon request from a peer, the central index will match the request with the best peer in its directory that matches the request. The best peer could be the one that is cheapest, fastest, or the most available, depending on the user needs. Then a file exchange will occur directly between the two peers. This model requires some managed infra-
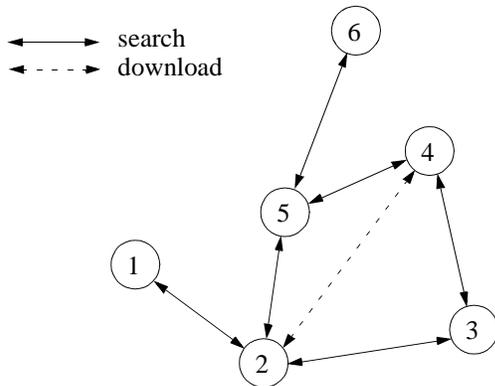
Figure 12: Flooded Requests Algorithm.

Figure 13: Document Routing Algorithm.

structure (the directory server), which hosts information about all participants in the community. This can cause the model to show some scalability limits, because it requires bigger servers when the number of requests increase, and larger storage when the number of users increase. However, Napster' experience showed that – except for legal issues – the model was very strong and efficient.

Flooded requests model. The flooding model is different from the central index one. This is a pure P2P model in which no advertisement of shared resources occurs. Instead, each request from a peer is *flooded* (broadcast) to directly connected peers, which themselves flood their peers etc., until the request is answered or a maximum number of flooding steps (typically 5 to 9) occur (see Figure 12). This model, which is used by Gnutella, requires a lot of network bandwidth, and hence does not prove to be very scalable, but it is efficient in limited communities such as a company network. To circumvent this problem, some companies have been developing "super-peer" client software, that concentrates lots of the requests. This leads to much lower network bandwidth requirement, at the expense of high CPU consumption. Caching of recent search requests is also used to improve scalability.

Document routing model. The document routing model, used by FreeNet, is the most recent approach. Each peer from the network is assigned a random ID and each peer also knows a given number of peers (see Figure 13). When a document is *published* (shared) on such a system, an ID is assigned to the document based on a hash of the document's contents and its name. Each peer will then *route* the document towards the peer with the ID that is most similar to the document ID. This process is repeated until the nearest peer ID is the current peer's ID. Each routing operation also ensures that a local copy of the document is kept. When a peer requests the document from the P2P system, the request will go to the peer with
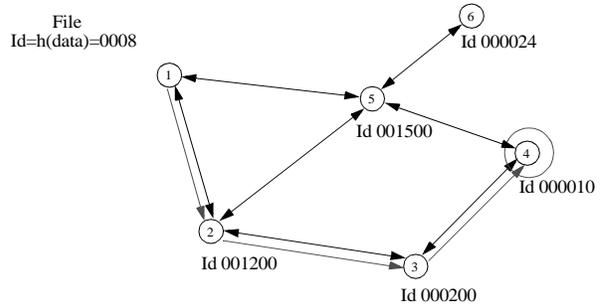
the ID most similar to the document ID. This process is repeated until a copy of the document is found. Then the document is transferred back to the request originator, while each peer participating the routing will keep a local copy.

Although the document routing model is very efficient for large, global communities, it has the problem that the document IDs must be known before posting a request for a given document. Hence it is more difficult to implement a search than in the flooded requests model. Also, network partitioning can lead to an *islanding* problem, where the community splits into independent sub-communities, that don't have links to each other.

Four main algorithms have implemented the document routing model: Chord, CAN, Tapestry, and Pastry. The goals of each algorithm are similar. The primary goals are to reduce the number of P2P hops that must be taken to locate a document of interest and to reduce the amount of routing state that must be kept at each peer. Each of the four algorithms either guarantee logarithmic bounds with respect to the size of the peer community, or argue that logarithmic bounds can be achieved with high probability. The differences in each approach are minimal, however each is more suitable for slightly different environments. In Chord, each peer keeps track of $\log N$ other peers (where N is the total number of peers in the community). When peer joins and leaves occur the highly optimized version of the algorithm will only need to notify $\log N$ other peers of the change. In CAN, each peer keeps track of only a small number of other peers (possibly less than $\log N$). Only this set of peers is affected during insertion and deletion, making CAN more suitable for dynamic communities. However, the tradeoff in this case lies in the fact that the smaller the routing table of a CAN peer, the longer the length of searches. Tapestry and Pastry are very similar. The primary benefit of these algorithms over the other two is that they actively try to reduce the latency of each P2P hop in addition to reducing the number of hops taken during a search.

P2P System	Algorithm Comparison Criteria					
	Model	**Parameters**	**Hops to locate data**	**Routing state**	**Peer joins and leaves**	**Reliability**
Napster	centralized metadata index location inquiry from central server download directly from peers	none	constant	constant	constant	Central server returns multiple download locations, client can retry
Gnutella	Broadcast request to as many peers as possible, download directly	none	no guarantee	constant (approx 3-7)	constant	receive multiple replies from peers with available data, requester can retry
Chord	uni-dimensional, circular ID space	N - number of peers in network	$\log N$	$\log N$	$(\log N)^2$	replicate data on multiple consecutive peers, app retries on failure
CAN	multidimensional ID space	N - number of peers in network d - number of dimensions	$d \cdot N^{1/d}$	$2 \cdot d$	$2 \cdot d$	multiple peers responsible for each data item, app retries on failure
Tapestry	Plaxton-style global mesh	N – number of peers in network b – base of the chosen identifier	$\log_b N$	$\log_b N$	$\log N$	replicate data across multiple peers, keep track of multiple paths to each peer
Pastry	Plaxton-style global mesh	N - number of peers in network b - base of the chosen identifier	$\log_b N$	$b \cdot \log_b N + b$	$\log N$	replicate data across multiple peers, keep track of multiple paths to each peer

Table 3. Comparison of Different P2P Location Algorithms.

Comparison of algorithms. The Chord algorithm models the identifier space as a uni-dimensional, circular identifier space. Peers are assigned IDs based on a hash on the IP address of the peer. When a peer joins the network, it contacts a gateway peer and routes toward its successor. The routing table at each peer n contains entries for $\log N$ other peers where the *i-th* peer succeeds n by at least 2^{i-1}. To route to another peer, the routing table at each hop is consulted and the message is forwarded toward the desired peer. When the successor of the new peer is found, the new peer takes responsibility for the set of documents that have identifiers less than or equal to its identifier and establishes its routing table. It then updates the routing state of all other peers in the network that are affected by the insertion. To increase the robustness of the algorithm, each document can be stored at some number of successive peers. Therefore, if a single peer fails, the network can be repaired and the document can be found at another peer.

CAN models the identifier space as multidimensional. Each peer keeps track of its neighbors in each dimension. When a new peer joins the network, it randomly chooses a point in the identifier space and contacts the peer currently responsible for that point. The contacted peer splits the entire space for which it is responsible into two pieces and transfers responsibility of half to the new peer. The new peer also contacts all of the neighbors to update their routing entries. To increase the robustness of this algorithm, the entire identifier space can be replicated to create two or more "realities". In each reality, each peer is responsible for a different set of information. Therefore, if a document cannot be found in one reality, a peer can use the routing information for a second reality to find the desired information.

Tapestry and Pastry are very similar and are based on the idea of a Plaxton mesh. Identifiers are assigned based on a hash on the IP address of each peer. When a peer joins the network, it contacts a gateway peer and routes toward the peer in the network with the ID that most closely matches the its own ID. Routing state for the new peer is built by copying the routing state of the peers along the path toward the new peer's location. For a given peer n, its routing table will contain *i* levels where the *i-th* level contains references to *b* nodes (where *b* is the base of the identifier) that have identifiers that match *n* in the last *i* positions. Routing is based on a longest suffix protocol that selects the next hop to be the peer that has a suffix that matches the desired location in the greatest number of positions. Robustness in this protocol relies on the fact that at each hop, multiple nodes, and hence multiple paths, may be traversed.

4 CHARACTERISTICS

This section addresses issues in P2P technology: decentralization, scalability, anonymity, self-organization, cost of ownership, ad-hoc connectivity, performance, security, transparency, usability, fault-resilience, and interoperability. These issues have a major impact on the effectiveness and deployment of P2P systems and applications. We compare them in the summary, Section 4.12.

4.1 Decentralization

P2P models question the wisdom of storing and processing data only on centralized servers and accessing the content via request-response protocols. In traditional client-server models, the information is concentrated in centrally located servers and distributed through networks to client computers that act primarily as user interface devices. Such centralized systems are ideal for some applications and tasks. For example, access rights and security are more easily managed in centralized systems. However, the topology of the centralized systems inevi-
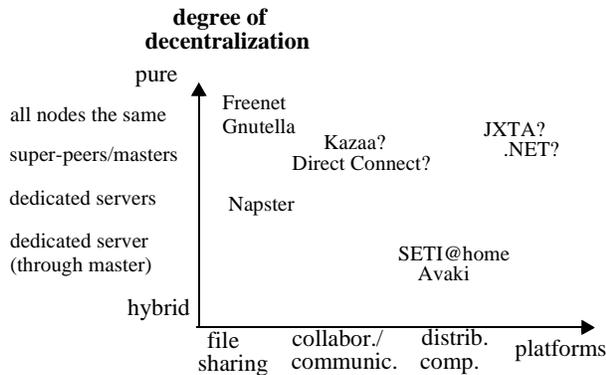
Figure 14: Examples of Levels of Decentralization in Various P2P Systems.

tably yields inefficiencies, bottlenecks, and wasted resources. Furthermore, although hardware performance and cost have improved, centralized repositories are expensive to set up and hard to maintain. They require human intelligence to build, and to keep the information they contain relevant and current.

One of the more powerful ideas of decentralization is the emphasis on the users' ownership and control of data and resources. In a fully decentralized system, every peer is an equal participant. This makes the implementation of the P2P models difficult in practice because there is no centralized server with a global view of all the peers in the network or the files they provide. This is the reason why many P2P file systems are built as hybrid approaches as in the case of Napster, where there is a centralized directory of the files but the nodes download files directly from their peers.

In fully decentralized file systems, such as Freenet and Gnutella, just finding the network becomes difficult. In Gnutella, for example, new nodes must know the address of another Gnutella node or use a host list with known IP addresses of other peers. The node joins the network of peers by establishing a connection with at least one peer currently in the network. Then, it can begin discovering other peers and cache their IP addresses locally.

One way to categorize the autonomy of a P2P system is through the "pure P2P" versus "hybrid P2P" distinction. A more precise decomposition may be as presented in Figure 14. This categorization has a direct effect on the self-organization and scalability of a system, as the purest systems are loosely coupled to any infrastructure.

4.2 Scalability

An immediate benefit of decentralization is improved scalability. Scalability is limited by factors such as the amount of centralized operations (e.g, synchronization and coordination) that needs to be performed, the amount of state that needs to be maintained, the inherent parallelism an application exhibits, and the programming model that is used to represent the computation.

Napster attacked the scalability problem by having the peers directly download music files from the peers that possess the requested document. As a result, Napster was able to scale up to over 6 million users at the peak of its service. In contrast, SETI@home [SETI@home 2001] focuses on a task that is embarrassingly parallel. It harnesses the computer power that is available over the Internet to analyze data collected from its telescopes with the goal of searching for extraterrestrial life forms. SETI@home has close to 3.5 million users so far. Systems like Avaki address scalability by providing a distributed object model.

Achieving good scalability should not be at the expense of other desirable features, such as determinism and performance guarantees. To address this problem, hybrid P2P systems, such as Napster, intentionally keep some amount of the operations and files centralized.

Early P2P systems such Gnutella [Gnutella 2001] and Freenet [Clark 2001] are ad-hoc in nature. A peer has to "blindly" send its requests to many other peers, causing the rest of the peers to search for the requested document. This can cause the time to retrieve a document to be unbounded. In addition, searching may fail even when an object exists, making the behavior of the system nondeterministic.

Recent P2P systems, represented by CAN, Chord, Oceanstore, and PAST, dictate a consistent mapping between an object key and hosting node. Therefore, an object can always be retrieved as long as the hosting nodes can be reached. Nodes in these systems compose an overlay network. Each node only maintains information about a small number of other nodes in the system. This limits the amount of state that needs to be maintained, and hence increases scalability. The logical topology of the overlay provides some guarantees on the lookup cost. These systems are designed to scale to billions of users, millions of servers and over 10^{14} files.

In the future, as the bandwidth and computation power continue to grow, platforms will be able to take advantage of this power, which should become interesting to more applications. The net effect is that these architectures will enable more automated scaling, as much resources can be provided, the applications could scale.

4.3 Anonymity

An important goal of anonymity is to allow people to use systems without concern for legal or other ramifications. A further goal is to guarantee that censorship of digital content is not possible. The authors of the Free Haven [Dingledine 2000] have identified the following forms of anonymity:

- Author: A document's author or creator cannot be identified
- Publisher: The person who published the document to the system cannot be identified
- Reader: People who read or otherwise consume data cannot be identified
- Server: Servers containing a document cannot be identified based on the document
- Document: Servers do not know what documents they are storing
- Query: A server cannot tell what document it is using to respond to a user's query

Regardless of the forms of anonymity, enforcing them typically involves enforcing three different kinds of anonymity between a communicating pair: *sender anonymity,* which hides the sender's identity; *receiver anonymity,* which hides a receiver's identity; and *mutual anonymity,* which hides the identities of the sender and receiver are hidden from each other and other peers [Pfitzmann 1987].

Besides the kinds of anonymity, it is also very important to understand the degree of anonymity a certain technique can achieve. Reiter and Rubin [Reiter 1998] presented a spectrum of anonymity degrees that cover *absolute privacy*, *beyond suspicion*, *probable innocence*, and *provably exposed*. For example, beyond suspicion means that even though an attacker can see evidence of a sent message., the sender appears no more likely to be the originator of that message than any other potential sender in the system.

There are six popular techniques, each suitable for enforcing different kinds of anonymity and with different kinds of constraints. We summarize them below.

Multicasting. Multicasting (or broadcasting) can be used to enforce receiver anonymity [Pfitzmann 1987]. A multicast group is formed for parties who wish to keep anonymous. An entity that is interested in obtaining a document subscribes to the multicast group. The party that possesses the document sends the document to the group. The identity of the requestor is effectively hidden from both the sender and other members of the group, and the requestor's anonymity is *beyond suspicion*. This technique can take advantage of the underlying network that supports multicast (e.g., Ethernet or token ring).

Spoofing the sender's address. For connectionless protocols such as UDP, the anonymity of the sender of a message can be enforced by spoofing the sender's IP address. This however, requires changing the protocol. In addition, this is not always feasible, because most ISPs now filer packets originating from invalid IP addresses.

Identity Spoofing: Besides changing the originator's address, anonymity can also be ensured by changing the identity of a communicating party. For example, in Freenet [Clark 2001], a peer passing a file to a requestor, either out of its own cache or from an upstream peer, can claim to be the owner of the content. The responder is *possibly innocent*, from an attackers's point view, because there is a nontrivial probability that the real responder is someone else.

Covert paths. Instead of communicating directly, two parties communicate through some middle nodes. Most existing techniques ensure only sender anonymity. A party that wishes to hide its identity prepares a covert path with the other party as the end of the path. Examples include Mix [Chaum 1981], Onion [Syverson 1997], Anonymizing Proxy [Gabber 1997], Crowds [Reiter 1998] and Herdes [Shields 2000]. The covert paths can use store/forward or persistent connection. By varying the length of the covert paths and changing the selected paths with different frequency, different degrees of anonymity can be achieved.

Intractable aliases. LPWA [Gabber 1999] is a proxy server that generates consistent untraceable aliases for clients from the servers. The client can open an account and be recognized upon returning to the opened account, while hiding the true identity of the client from the server. Techniques of this kind ensure sender anonymity and rely on a trusted proxy server. The degree of anonymity that can be achieved falls in between *absolute privacy* and *beyond suspicion*.

Non-voluntary placement. An interesting new approach is anonymity via non-voluntary placement of a document on a hosting node. Here, a publisher forces a document onto a hosting node using, for example, consistent hashing. Because the placement is non-voluntary, the host cannot be held accountable for owning the document.

Project	Types and Techniques of Anonymity			
	Publisher	Reader	Server	Document
Gnutella	multicasting, covert paths	N/A	N/A	N/A
Freenet	covert path, identity spoofing	covert paths	non-voluntary placement	encryption
APFS	covert paths	covert paths	N/A	N/A
Free Haven	covert paths (remailer)	covert paths	broadcast	encryption/split files into shares
Publius	covert paths (remailer)	N/A	non-voluntary placement	encryption/ split key
PAST	N/A	N/A	non-voluntary placement	encryption

Table 4. Types of Anonymity and Techniques to Enforce Them

We now summarize the forms of anonymity some of the popular P2P systems support and the techniques they employ (see Table 4).

Gnutella [Gnutella 2001] and Freenet [Clark 2001] provide a certain degree of anonymity in the way peers request/send a document. In Gnutella, a request is broadcast and rebroadcast until it reaches a peer with the content. In Freenet, a request is sent and forwarded to a peer that is most likely to have the content. The reply is sent back along the same path.

APFS [Scarlata 2001] addresses the mutual anonymity problem assuming that trusted centralized support does not exist. Peers may inform a (untrusted) coordinator about their willingness to be index servers. Both the initiator and the responder need to prepare their own covert paths.

Free Haven [Dingledine 2000] and Publius [Waldman 2000] are designed to defend against censorship. They strengthen the document's anonymity by further splitting the files as they are stored at each server. In this way, no single server even contains all of the data needed to perform an attack on the encrypted file contents. The anonymity among a pair of communicating parties (publisher/server, reader/server) is enforced via *covert paths*. Both Free Haven and Publius build the covert paths using an anonymous re-mailer system. Publius could be extended to support reader anonymity by using its re-mailer for publishing, but it does not currently do so.

PAST [Rowstron 2001], CAN [Ratnasamy 2001] and Chord [Stoica 2001] represent a new class of P2P system that provides a reliable infrastructure. One common property among these systems is that object placement can be entirely non-voluntary. As a result, when an object is placed on a node, that node cannot be held accountable for owning that object. The embedded routing mechanisms in these systems can also easily be adapted to covert path for mutual anonymity.

4.4 Self-Organization

In cybernetics, self-organization is defined as "a process where the organization (constraint, redundancy) of a system spontaneously increases, i.e., without this increase being controlled by the environment or an encompassing or otherwise external system" [Heylighen 1997].

In P2P systems, self-organization is needed because of scalability, fault resilience, intermittent connection of resources, and the cost of ownership. P2P systems can scale unpredictably in terms of the number of systems, number of users, and the load. It is very hard to predict any one of them, requiring frequent re-configuration of centralized system. The significant level of scale results in an increased probability of failures, which requires self-maintenance and self-repair of the systems. Similar reasoning applies to intermittent disconnection; it is hard for any predefined configuration to remain intact over a long period of time. Adaptation is required to handle the changes caused by peers connecting and disconnecting from the P2P systems. Finally, because it would be costly to have dedicated equipment and/or people for managing such a fluctuating environment, the management is distributed among the peers.

There are a number of academic systems and products that address self-organization. In OceanStore, self-organization is applied to location and routing infrastructure [Kubiatowicz et al 2000, Rhea et al. 2001, Zhao et al. 2001]. Because of intermittent peer availability, as well as variances in network latency and bandwidth, the infrastructure is continuously adapting its routing and location support.

In Pastry, self-organization is handled through protocols for node arrivals and departures based on a fault-tolerant overlay network [Druschel and Rowstron 2001, Rowstron and Druschel 2001]. Client requests are guaranteed to be routed in less than $\lceil \log_{16} N \rceil$ steps on average. Also, file replicas are distributed and storage is randomizes for load balancing.

The FastTrack product attributes quicker search and download to self-organizing distributed networks. In these networks, more powerful computers automatically become SuperNodes and act as search hubs. Any client can become a SuperNode if it meets processing, and networking criteria (bandwidth and latency). The distributed networks replace any centralized service [FastTrack 2001].

SearchLing uses self-organization to adapt its network according to the type of search, resulting in reduced network traffic and less unreachable information [SearchLing 2001].

4.5 Cost of Ownership

One of the premises of P2P computing is shared ownership. Shared ownership reduces the cost of owning the systems and the content, and the cost of maintaining them. This is applicable to all classes of P2P systems. It is probably most obvious in distributed computing. For example, SETI@home is faster than the fastest supercomputer in the world, yet at only a fraction of its cost – 1% [Anderson 2000].

The whole concept of Napster music sharing was based on each member contributing to the pool of music files. Similar assumptions for peers are used in other file systems, such as OceanStore.

In P2P collaboration and communication systems, and in platforms, elimination of centralized computers for storing information also provides reduced ownership and maintenance costs. A similar approach is taken in wireless communication in the United States. A so-called "Parasitic grid" wireless movement, enables sharing of the existing home-installed 802.11b bandwidth among the users [Schwarz 2001]. These networks compete with the companies installing wireless infrastructure at the fraction of the cost.

4.6 Ad-Hoc Connectivity

The ad-hoc nature of connectivity has a strong effect on all classes of P2P systems. In distributed computing, the parallelized applications cannot be executed on all systems all of the time; some of the systems will be available all of the time, some will be available part of the time, and some will be not be available at all. P2P systems and applications in distributed computing need to be aware of this ad-hoc nature and be able to handle systems joining and withdrawing from the pool of available P2P systems. While in traditional distributed systems, this was an exceptional event, in P2P systems it is considered usual.

In content sharing P2P systems and applications, users expect to be able to access content intermittently, subject to the connectivity of the content providers. In systems with higher guarantees, such as service-level agreements, the ad-hoc nature is reduced by redundant service providers, but the parts of the providers may still be unavailable.

In collaborative P2P systems and applications, the ad-hoc nature of connectivity is even more evident. Collaborative users are increasingly expected to use mobile devices, making them more connected to Internet and available for collaboration. To handle this situation, collaborative systems support transparent delay of communication to disconnected systems. This can be accomplished by having proxies delegated on networks to receive messages, or by having other sorts of relays on the sending system or somewhere in the network that will temporarily hold communication for an unavailable system.

Furthermore, not everything will be connected to the Internet. Even under these circumstance, ad-hoc groups of people should be able to form ad-hoc networks in order to collaborate. The supporting ad-hoc networking infrastructures, such as 802.11b, Bluetooth, and infrared, have only a limited radius of accessibility. Therefore, both P2P systems and applications need to be designed to tolerate sudden disconnection and ad-hoc additions to groups of peers.

4.7 Performance

Performance is a significant concern in P2P systems. P2P systems aim to improve performance by aggregating distributed storage capacity (e.g., Napster, Gnutella) and computing cycles (e.g., SETI@Home) of devices spread across a network. Because of the decentralized nature of these models, performance is influenced by three types of resources: processing, storage, and networking. In particular, networking delays can be significant in wide-area networks. Bandwidth is a major factor when a large number of messages are propagated in the network and large amounts of files are being transferred among many peers. This limits the scalability of the system. Performance in this context does not put emphasis in the millisecond level, but rather tries to answer questions of how

long it takes to retrieve a file or how much bandwidth will a query consume.

In centrally coordinated systems (e.g., Napster, Seti@Home) coordination between peers is controlled and mediated by a central server, although the peers also may later contact each other directly. This makes these systems vulnerable to the problems facing centralized servers. To overcome the limitations of a centralized coordinator, different hybrid P2P architectures [Yang, 2001] have been proposed to distribute the functionality of the coordinator in multiple indexing servers that cooperate with each other to satisfy user requests. DNS is another example of a hierarchical P2P system that improves performance by defining a tree of coordinators, with each coordinator responsible for a peer group. Communication between peers in different groups is achieved through a higher level coordinator.

In decentralized coordinated systems such as Gnutella and Freenet, there is no central coordinator; communication is handled individually by each peer. Typically, they use message forwarding mechanisms search for information and data. The problem with such systems is that they end up sending a large number of messages over many hops from one peer to another. Each hop contributes to an increase in the bandwidth on the communication links and to the time required to get results for the queries. The bandwidth for a search query is proportional to the number of messages sent, which in turn is proportional to the number of peers that must process the request before finding the data.

There are three key approaches to optimize performance: replication, caching, and intelligent routing.

Replication. Replication puts copies of objects/files closer to the requesting peers, thus minimizing the connection distance between the peers requesting and providing the objects. Changes to data objects have to be propagated to all the object replicas. Oceanstore uses an update propagation scheme based on conflict resolution that supports a wide range of consistency semantics. The geographic distribution of the peers helps to reduce congestion on both the peers and the network. In combination with intelligent routing, replication helps to minimize the distance delay by sending requests to closely located peers. Replication also helps to cope with the disappearance of peers. Because peers tend to be user machines rather than dedicated servers, there is no guarantee that the peers won't be disconnected from the network at random.

Caching. Caching reduces the path length required to fetch a file/object and therefore the number of messages exchanged between the peers. Reducing such transmissions is important because the communication latency between the peers is a serious performance bottleneck facing P2P systems. In Freenet for example, when a file is found and propagated to the requesting node, the file is cached locally in all the nodes in the return path. More efficient caching strategies can be used to cache large amounts of data infrequently. The goal of caching is to minimize peer access latencies, to maximize query throughput and to balance the workload in the system. The object replicas can be used for load balancing and latency reduction.

Intelligent routing and network organization. To fully realize the potential of P2P networks, it is important to understand and explore the social interactions between the peers. The most pioneering work in studying the social connections among people is the "small-world phenomenon" initiated by Milgram [Milgram 1967]. The goal of his experiment was to find short chains of acquaintances linking pairs of people in the United States who did not know one another. Using booklets of postcards he discovered that Americans in the 1960s were, on average, about six acquaintances away from each other. Adamic et al. have explored the power-law distribution of the P2P networks, and have introduced local search strategies that use high-degree nodes and have costs that scale sub-linearly with the size of the network [Adamic et al. 2001]. Ramanathan et al [Ramanathan, 2001] determine "good" peers based on interests, and dynamically manipulate the connections between peers to guarantee that peers with a high degree of similar interests are connected closely. Establishing a good set of peers reduces the number of messages broadcast in the network and the number of peers that process a request before a result is found. A number of academic systems (Oceanstore, Pastry, see Section 2.6)) improve performance by proactively moving the data in the network. The advantage of these approaches is that peers decide whom to contact and when to add/drop a connection based on local information only.

4.8 Security

P2P systems share most of their security needs with common distributed systems: trust chains between peers and shared objects, session key exchange schemes, encryption, digital digests, and signatures. Extensive research has been done in these areas, and we will not discuss it further in the present document. New security requirements appeared with P2P systems.

- **Multi-key encryption.** File sharing systems such as Publius intend to protect a shared object, as well as the anonymity of its author, publishing peer and hosting peer. The security scheme chosen by Publius developers is based on a (Public key, Multiple private keys) asymmetric encryption mechanism derived from R. Shamir's "shared secrets" encryption method [Shamir 1979]. Byzantine attacks by malicious authenticated users have typically been an issue for such schemes. Recent improvements (see [Castro and Liskov 2001] for an example) have greatly reduced the costs inherent to Byzantine agreements and opened the way to solid systems used by large numbers of users.

- **Sandboxing.** Distributed computing P2P systems require execution of some code on peer machines. It is crucial to protect the peer machines from potentially malicious code and protect the code from a malicious peer machine. Protecting a peer machine typically involves enforcing (1) safety properties such that the external code will not crash the host box, or will only access the host data in a type-safe way, and (2) enforcing security properties to prevent sensitive data from being leaked to malicious parties. Techniques to enforce the first include sandboxing, safe languages (e.g., java), virtual machines (e.g., Internet C++ POSIX virtual machine, real mode Linux derivatives, which run a virtual machine on top of the actual OS, VMware), proof-carrying code and certifying compilers [Necula 1997, 1998] and program verification techniques applied to verifying the safety properties of machine code [Xu 2000, 2001]. Techniques to check the latter include information flow theory [Denning 1976], and model checking [Ball 2001].

- **Digital Rights Management.** P2P file sharing makes file copying easy. It is necessary to be able to protect the authors from having their intellectual property stolen. One way to handle this problem is to add a signature *in* the file that makes it recognizable (the signature remains attached to the file contents) although the file contents do not seem affected. This technique, referenced as *watermarking* or *steganography* [Katzenbeisser 1999], has been experimented with by RIAA to protect audio files such as MP3s, hiding the Copyright information in the file in inaudible ways.

- **Reputation and Accountability.** We already spoke about *trust*, which is the way we will mathematically ensure that a communiquee is actually the entity it claims it is. In P2P systems, reputation is built on top of trust, and requires ways to measure how "good" or "useful" a peer is. For instance, if a given user shares lots of interesting files, its reputation should be high. *Freeloader* is a common term for a user who downloads files from P2P systems without offering files to others. A freeloader usually has a low reputation. To prevent this kind of non-cooperative behavior, some accountability mechanisms need to be devised. Current systems often rely on cross-ratings, but because it is based on a community of *authenticated* but *untrusted* users, it is difficult to produce a solid system.

- **Firewalls.** P2P applications inherently require direct connections between peers. However, in corporate environments internal networks get isolated from the external network (the Internet), leaving reduced access rights to applications. For instance, most firewalls block inbound TCP connections [Peer-to-peer Working Group. 2001]. This means that a machine within a Firewall will not be accessible from a machine external to the network. Worse, home users frequently use IP Masquerading or Network Address Translation (NAT) technology to share an internet connection between several machines, which leads to the same inaccessibility problem. However, as outbound access through port 80 (HTTP) is often allowed by firewalls, some mechanisms have been devised that enable connections between hidden (machines behind a firewall or NAT, inaccessible from the Internet) and Internet machines. This is quite limiting however, as it requires connection to be initiated from the hidden machine. When both peers who want to communicate reside behind different firewalls, the problem becomes harder. It requires a central reflector (or relay) server on the Internet, which provides a connection between the hidden peers.

4.9 Transparency and Usability

In distributed systems, transparency was traditionally associated with the ability to transparently connect distributed systems into a seamlessly local system. The primary form of transparency was *location transparency*, but other forms included transparency of *access*, *concurrency*, *replication*, *failure*, *mobility*, *scaling*, etc. [Coulouris1990]. Over time, some of the transparencies were further qualified, such as transparency for failure, by requiring distributed applications to be aware of failures [Waldo et al 1997], and addressing transparency on the Internet and Web (see next paragraph).

From its beginning, the Internet paid particular attention to transparency at the protocol level (TCP/IP), so called *end-to-end address transparency* [Carpenter 2000]. The end-to-end argument [Saltzer et al. 1984] claims that certain functions in a communication between two entities can be performed only with the knowledge and state maintained at these entities at the application level (hence, end-to-end: all the way from application through the communication stack, up to the other application).

18

This implies that application communication state is not maintained in the network, but rather at the communication end points.

This also implies that any point in the network knows the name and address of the other communication point, an assumption that is not true any more. IPv4's lack of the domain names and IP numbers, as well as the introduction of intranets and mobile users, resulted in IP numbers that are valid only during a single session. Examples include SLIP and PPP, VPNs, use of firewalls, DHCP, private addresses, network address translators (NATs), split DNS, load sharing optimizations, etc [Carpenter 2000]. This had significant implications for P2P and was also one of the reasons for the introduction of P2P. Because it was no longer possible to rely on DNS to provide an accurate name, P2P systems came up with different naming and discovery schemes (see Section 3.1 as well as end of Section 5.3).

Web naming did not necessarily offer full naming transparency. URLs are widely used instead of URNs, which were supposed to enable *naming transparency* [Berners-Lee et al. 1998]. Beside *naming/addressing transparency* in P2P there is also a requirement for *administration transparency*. Users are typically non-experts and they do not or cannot administer their software and devices.

The P2P software should not require any significant set up or configuration of either networks or devices in order to be able to run. Also, self-updating software is a desirable feature. In addition, P2P systems should be network and device transparent (independent). They should work on the Internet, intranets, and private networks, using high-speed or dial-up links. They should also be device transparent, which means they should work on a variety of devices, such as handheld personal digital assistants (PDAs), desktops, cell phones, and tablets.

Another form of transparency is related to security and mobility. Automatic and transparent authentication of users and delegation to user proxies can significantly simplify users' actions. Supporting mobile users and disconnection in particular, can enable users to work independently of whether and how they are connected to the Internet or intranets.

A user can use P2P applications in the following manners:

- as a user of services, typically through Web interfaces (e.g., content sharing, information gathering)
- wrapped around non-P2P applications, typically on a P2P platform (e.g., Groove, .NET)
- as locally installed P2P software (e.g., distributed computing screensavers and Napster)

4.10 Fault Resilience

One of the primary design goals of a P2P system is to avoid a central point of failure. Although most P2P systems (pure P2P) already do this, they nevertheless are faced with failures commonly associated with systems spanning multiple hosts and networks: disconnections/ unreachability, partitions, and node failures. These failures may be more pronounced in some networks (e.g., wireless) than others (e.g., wired enterprise networks). It would be desirable to continue active collaboration among the still connected peers in the presence of such failures. An example would be an application, such as genome@home [Genome@HOME 2001] executing a partitioned computation among connected peers. Would it be possible to continue the computation if one of the peers were to disappear because of a network link failure? If the disconnected peer were to reappear, could the completed results (generated during the standalone phase) be integrated into the ongoing computation? Questions similar to these would have to be addressed by P2P systems aiming to provide more than just "best effort" Internet service.

In the past, client-server disconnection has been studied for distributed file systems that consider mobile clients (e.g., Coda [Satyanarayanan 1990]), and a common solution is to have application-specific resolvers to handle any inconsistency on reconnection. Some current P2P systems (e.g., Groove [Groove 2001]) handle this by providing special nodes, called relays, that store any updates or communication temporarily until the destination (in this case another peer) reappears on the network. Others (e.g., Magi [Magi 2001]) queue messages at the source, until the presence of the destination peer is detected.

Another problem related to disconnection is non-availability of resources. This may occur either because the resource is unreachable because of a network failure or because the peer hosting the resource has crashed/gone offline. While the former may be resolved by routing around the failure and is already supported by the Internet, the latter requires more careful consideration. Replication of crucial resources helps alleviate the problem. P2P networks such as Napster and Gnutella represent systems having both a passive and an uncontrolled replication mechanism based solely on the file's popularity. Depending on the application running over these networks, it may be necessary to provide certain persistence guarantees. This requires a more active and reliable replication policy.

Anonymous publishing systems such as Freenet [Clark et al. 2000] and Publius [Waldman 2000] ensure avail-

ability by controlled replication. Oceanstore [Kubiatow-icz 2000] maintains a two-layered hierarchy of replicas and through monitoring of administrative domains avoids sending replicas to locations with highly correlated probability of failures. However, because a resource in the P2P system could be more than a just a file – such as a proxy to the Internet, shared storage space, or shared computing power – the concepts of replicated file systems have to be extended to additional types of resources. Grid computing solutions (e.g. Legion) provide resilience against node failures by restarting computations on different nodes.

A challenging aspect of P2P systems is that the system maintenance responsibility is completely distributed and needs to be addressed by each peer to ensure availability. This is quite different from client-server systems, where availability is a server-side responsibility.

4.11 Interoperability

Although many P2P systems already exist there is still no support to enable these P2P systems to interoperate. Some of the requirements for interoperability include:

- How do systems determine that they can ineroperate
- How do systems communicate, e.g., what protocol should be used, such as sockets, messages, or HTTP
- How do systems exchange requests and data, and execute tasks at the higher level, e.g., do they exchange files or search for data
- How do systems determine if they are compatible at the higher protocol levels, e.g., can one system rely on another to properly search for a piece of information
- How do systems advertise and maintain the same level of security, QoS, and reliability

In the past, there were different ways to approach interoperability, such as *standards* IEEE, (e.g., IEEE standards for ethernet, token ring, and wireless); *common specifications*, (e.g., Object Management Group [OMG, 2001]); *common source code*, (e.g., OSF DCE [Rosenberry, 1992]); *open-source* (e.g., Linux); and *de facto standards* (e.g., Windows or Java).

In the P2P space, some efforts have been made towards improved interoperability, even though interoperability is still not supported. The P2P Working Group [p2pwg, 2001] is an attempt to gather the community of P2P developers together and establish common ground by writing reports and white papers that would enable common understanding among P2P developers. The P2P Working Group gathers developers from both ad-hoc communication systems and grid systems. The Grid Forum is a sim-

ilar effort in the grid computing space. Both efforts represent an approach similar to OMG, in defining specifications and possibly reference implementations.

The JXTA effort [JXTA, 2001] approaches interoperability as an open-source effort, by attempting to impose a de facto standard. A number of developers are invited to contribute to the common source tree with different pieces of functionality. Only a minimal underlying architecture is supported as a base, enabling other systems to contribute parts that may be compatible with their own implementations. A number of existing P2P systems have already been ported to the JXTA base. JXTA is described in more detail in Section 6.7.

4.12 Summary

Decentralization is a key feature of P2P systems. It affects how developers design systems and applications, by influencing algorithms, data structures, security, scalability, and availability assumptions. It affects how users can be connected to a system and the people with whom they can interact. For example, in games, users perceive that other players are remote, and that they can also be disconnected. This implies that they should devise strategies in a decentralized fashion. Distributed P2P applications are written assuming decentralization, and collaborative applications have to handle group management without central naming, authorization, and data repositories.

The ad-hoc nature of P2P systems also affects the way applications and systems are conceived. The fact that any system or user can disappear at time drives the design of these systems as well as user perceptions and expectations. In addition to the classical security issues of traditional distributed systems, P2P is distinguished by the importance of anonymity in certain applications and markets. Scalability, performance, fault resilience, and interoperability have similar importance for P2P as they have in traditional distributed systems.

Distributed computing applications are primarily concerned with scalability (which is derived from decentralization) and performance. Fault resilience is tied in with the ad-hoc nature of connectivity – distributed application designers and users need to account for the possibility of a system going down at any time and being able to resume from a previous checkpoint. P2P systems that use critical or confidential data are concerned with security. Content sharing applications and systems are primarily concerned with the availability of data. Enterprise systems are concerned with security and performance, and public and consumer systems are concerned with transparency (ease of use) and anonymity. Collaborative and

communication applications are concerned with connectivity (ad-hoc nature), security, and anonymity, and with interoperability between systems.

5 P2P SYSTEMS

This section describes in more detail the four categories presented in Figure 6: distributed computing, file sharing, collaborative systems, and P2P platforms. In addition, we also present historical P2P systems that predated the recent notion of P2P systems. While this section presents the P2P systems categories in general, Section 6 presents each of the categories in more detail with two case studies. This section and Section 6 are used as the basis for a comparison of systems in Appendix A, and specifically for Table 6.

5.1 Historical

In most cases, early distributed applications were P2P. When most users were from a technical or academic background, and were using either time-sharing systems or engineering workstations, P2P applications seemed the most natural approach. It was the late-80's and early-90's when client-server architectures became more prevalent because they provided the fastest and most cost-effective means of supporting large numbers of non-technical users. It also allowed the use of less expensive and less functional computers such as desktop PCs.

While most early distributed applications can be considered P2P, e-mail systems built on the Simple Mail Transfer Protocol (SMTP) and Usenet News were probably the most widely used. In each case, local servers that received a message built connections with peer servers to deliver messages into a user's mail file or into a spool file containing messages for the newsgroup. The File Transfer Protocol (FTP) was the precursor to today's file sharing P2P systems. While FTP is a client-server application, it was very common for individuals to run FTP servers on their workstations to provide files to their peers. Eventually, an indexing system, Archie, was developed to provide a central search mechanism for files on FTP servers. This structure with central search and distributed files is exactly replicated in the Napster P2P system.

Prior to the establishment of a continuously connected network such as the Internet, decentralized dial-up networks were widely used. The most notable examples include UUNet and Fidonet. These networks were composed of a collection of machines that made periodic dial-up connections to one another. On a typical connection, messages (again, typically e-mail or discussion group entries) were transferred bi-directionally. Often, a message would be routed through multiple dial-up hops to reach its destination. This multi-hop message routing approach can be seen in current P2P systems such as Gnutella.

In our "modern" era dominated by PCs on workplace LANs or home dial-up connections, the first widen use of P2P seems to have been in instant messaging systems such as AOL Instant Messenger. These are typically hybrid P2P solutions with discovery and brokering handled by a central server followed by direct communication between the peer messaging systems on the PCs. The current phase of interest and activity in P2P was driven by the introduction of Napster [Napster 2001] in 1999. It came at a time when computers and their network connections were nearing the level found previously in technical and academic environments, and re-created earlier approaches with an interface more suitable to a non-technical audience.

5.2 Distributed Computing

Distributed computing is very successfully used by P2P systems. The idea of using spare computing resources has been addressed for some time. The *Beowulf* project from NASA [Becker *et al.* 1995] was a major milestone that showed that high performance can be obtained by using a number of standard machines. Other efforts such as *MOSIX* [Barak and Litman 1985, Barak and Wheeler 1989] and *Condor* [Litzkow et al 1988, Litzkow and Solomon 1992] also addressed distributed computing in a community of machines, focusing on the *delegation* or *migration* of computing tasks from machine to machine.

Grid computing is another concept that was first explored in the 1995 I-WAY experiment [I. Foster], in which high-speed networks were used to connect high-end resources at 17 sites across North America. Out of this activity grew a number of Grid research projects that developed the core technologies for "production" Grids in various communities and scientific disciplines. Grid technology efforts are now focused around the Global Grid forum (http://www.globalgridforum.org) and the Globus project (http://www.globus.org). A computing grid can be seen and used as a single, transparent computer. A user logs in, starts jobs, moves files, and receives results in a standard way.

Derivatives of Grid Computing based on collaboration of standard Internet-connected PCs began to appear in the late 90's. In this document, we will use "*Distributed Computing*" terminology to describe them.

Distributed Computing achieves processing scalability by aggregating the resources of large number of individual Internet PCs. Typically, distributed computing requires applications that are run in a proprietary way by a central controller. Such applications are usually targeting massive multi-parameters systems, with long running jobs (months or years) using P2P foundations. One of the first widely visible distributed computing events occurred in January 1999, where distributed.net, with the help of several tens of thousands of Internet computers, broke the RSA challenge [DES-III] in less than 24 hours using a distributed computing approach. This made people realize how much power can be available from idle Internet PCs.

More recent projects have been raising interest from many users within the Internet community. For example, SETI@home [SETI@home 2001] now has a consolidated power of about 25 Tflop/s (Thousands of Billions of floating point operation per second), collected from more than three million registered user machines.

Peer-to-Peer? A common misconception is that distributed computing systems such as SETI@home are not P2P systems: The argument is that central server is required for controlling the offered PC resources, the PCs do not operate as servers, and no communication occurs between peers (PCs). However, a very significant part of the system is executed on the PCs, with high autonomy. Hence, we consider distributed computing systems to be P2P systems.

How it works. The computational problem to be solved is split into small independent parts. The processing of each of the parts (using a *fork-and-join* mechanism) is done on an individual PC and the results are collected on a central server. This central server is responsible for distributing job items to PCs on the Internet. Each of the registered PCs is equipped with client software. The client software takes advantage of inactivity periods (often characterized by screensaver activation times) in the PC to perform some computation requested by the server. After the computation is finished, the result is sent back to the server, and a new job is allocated to the client.

Current usage. One of the major limitations of Desktop Computing is that it requires jobs that can be split into *independent* small parts that do not require cross-peer communication. Internet latencies do not allow demanding communication patterns such as those found in typical cluster computing. Hence, it is not possible to execute supercomputing-like processing such as linear algebra problems (matrix computation). Current applications
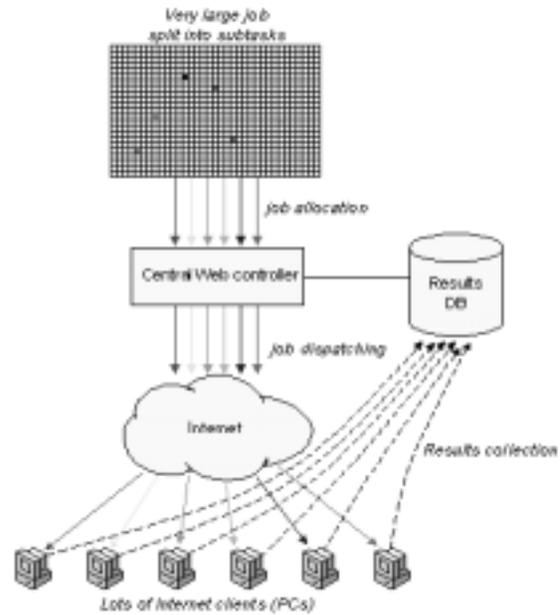

Figure 15: Distributed computing over the Web

consist of Single Process Multiple Data (SPMD) problems and multiprogramming problems where a given job has to be run on many different input data sets. This mostly includes simulation and model validation tasks. Because specific applications have to be developed using very specific constraints, paradigms, and environments, their development cost is prohibitive to most users, and the scope remains limited to highly visible research domains, such as the human genome project, alien seeking, cancer research, and weather model validation.

Application area examples – Financial and Biotechnology. Financial and biotechnology applications are suitable for distributed computing. Financial institutions, such as banks and credit companies, are executing complex simulations for market calculations. Applications include portfolio pricing, risk hedge calculation, market and credit evaluation, counter-party netting, margin calculations, cash flow, and demographic analysis [Intel 2001]. In the past, financial applications were typically run during the night. As they become more real-time in nature, these requirements will grow even further. So far only big institutions have been able to afford the computing power to automate these simulations, typically using mainframe computers or very powerful workstations. By relying on P2P commodity platforms, smaller banks can also benefit from these applications. Furthermore, as technology favors farms of desktop computers, they become not only more cost effective, but also a more powerful platform. As an example, Intel and DataSynapse claim speed-ups from 15 hours to 30 minutes in the case

of interest rate swap modeling when moving to a P2P solution [DataSynapse 2001]. The biggest challenge in using P2P for financial markets is the intrinsic requirement for security. Because most of these applications are executing behind the firewall, security requirements are somewhat relaxed. In particular, parts of the applications can even be executed outside of the firewall.

In the biotechnology sector, the need for advanced computing techniques is being driven by the availability of colossal amounts of data. For instance, genomic research has close to three billion sequences in the human genome database. Applying statistical inference techniques to data of this magnitude requires unprecedented computational power. Traditionally, scientists have used high-performance clustering (HPC) and supercomputing solutions, and have been forced to employ approximating techniques in order to complete studies in an acceptable amount of time. By harnessing idle computing cycles (95%-98% unused) from general purpose machines on the network, and grouping multi-site resources, grid computing makes more computing power available to researchers. Grid solutions partition the problem space among the aggregated resources to speed up completion times. Companies such as Platform Computing (LSF) [Platform Computing 2001], Entropia [Entropia], Avaki [Avaki 2001] and Grid Computing Bioinformatics [GCB 2001] offer complete HPC and grid solutions to biological research organizations and pharmaceutical research and development. Genomics and proteomics projects such as Genome@home and Folding@home managed by groups at Stanford make use of the idle cycles of registered clients to compute parts of the complex genome sequencing and protein folding problems.

In search for a business model. Distributed computing requires an expensive infrastructure to host the Web service, and this does not even account for development and maintenance costs. Attempts have been made by several companies (such as Porivo) to pay the user for use of his machine by the system. If your machine computes a lot of jobs, you can get some money for having shared your machine. Other models are based on auctions, where the cheapest available CPU is allocated for a given job. Another model used by Entropia addresses vertical markets, such as Genomics, which is based on huge database look ups. However, none of the existing business models have yet proven successful. The only model that currently seems to work is the dedicated community model, such as astronomy researchers sharing their machines in a "friendly" way.

5.3 File Sharing

Content storage and exchange is one of the areas where P2P technology has been most successful. Multimedia content, for instance, inherently requires large files. Napster and Gnutella have been used by Internet users to circumvent bandwidth limitations that make large file transfers unacceptable with classic mechanisms. Distributed storage systems based on P2P technologies are taking advantage of the existing infrastructure to offer the following features.

- *File exchange areas.* Systems such as Freenet provide the user with a potentially unlimited storage area by taking advantage of redundancy. A given file is stored on some nodes in the P2P community, but it is made available to any of the peers. A peer requesting a given file just has to know a reference to a file, and is able to retrieve the file from the community by submitting the file reference. Systems such as Freenet, Gnutella, and Kazaa fall in this category.

- *Highly available safe storage.* The duplication and redundancy policies in some projects offer virtual storage places where critical files get replicated multiple times, which helps ensuring their availability. *OceanStore* [Kubiatowicz et al. 2000] and *Chord* [Dabek et al. 2000] are examples of such systems.

- *Anonymity.* Some P2P systems such as Publius [Waldman, Rubin and Cranor references] mathematically ensure that published documents preserve anonymity for their authors and publishers, while allowing people to access the documents.

- *Manageability.* P2P systems enable easy and fast retrieval of the data by distributing the data to caches located at the edges of the network. The location of the data is not known by the retriever, perhaps not even after the data is retrieved. Freenet, for example, stores the data in many locations in the path between the provider and the retriever, so the whole notion of hosting a file becomes meaningless. Files move freely among the peers and are allowed to disappear even if they are being downloaded. This has some important implications. For example, the question is who is accountable for the files (see Section 4.5). Also, how can we ensure that the entire piece of data is being downloaded and cope with the un-reliability of the peers (see Section 4.10).

Technical issues. The major technical issues in file sharing systems are mostly the network bandwidth consumption, security, and search capabilities. Three main models exist today, as discussed in Section 3.2: *Centralized directory* model (Napster), the *flooded request* model (Gnutella), and the *document routing* model

(FreeNet). All P2P file sharing systems can be categorized into one of these three families, although variations do exist, such as extensions of the models with *leader election mechanisms* (automatic or manual elections) such as in *KaZaA,* which allow for better scalability of the systems and less stress on the network. In addition, current P2P systems (e.g., Napster, Gnutella, Freenet) have mainly focused on the exchange and sharing of "small" objects such as files and music clips. However, we expect that in future P2P systems the content will be of any form, including audio, video, software, and documents. To do that, we will need intelligent decisions such as from where the content be retrieved and over which network path should the content travel. XDegrees [XDegrees 2001], for example, ensures that information is efficiently routed to users and that multiple document replicas are synchronized across many peers. They provide an eXtensible Resource Name System (XRNS) based on the notion of the URL that creates a location-independent name space. They place frequently accessed information at optimal locations in the network and then select the best route to retrieve that information based on source availability, network topology, and response time.

Application area example. Napster is first P2P file sharing application that jump started the P2P area. Napster was originally developed to defeat the copying problem and to enable the sharing of music files over the Internet. Napster uses the centralized directory model (see Section 2.6) to maintain a list of music files, where the files are added and removed as individual users connect and disconnect from the system. Users submit search requests based on keywords such as "title," "artist," etc. Although Napster's search mechanism is centralized, the file sharing mechanism is decentralized; the actual transfer of files is done directly between the peers. Napster's centralized directory model inevitably yields scalability limitations (see Section 3.2). For example, your available bandwidth can be tremendously reduced by users downloading songs from your machine. Yet, centralization simplifies the problem of obtaining a namespace and enables the realization of security mechanisms (see Section 3.8).

Napster has been quite popular. It has had more than forty million client downloads and has led to numerous variants of file-sharing applications (such as Gnutella and Pointera). OpenNap (http://opennap.sourceforge.net/) is an open-source Napster server that extends the Napster protocol to allow sharing of any media type and add the ability to link Napster servers together. To solve the problem with copyright violations, Napster relaunched a new Napster membership service to start a new chapter in the music file business. Although it still offered the core functions (searching, finding, sharing, and discovering digital music through the Napster community) it also offered to artists the opportunity to register as rights holders and get paid for sharing their music on Napster. Napster set the rules for how their music files are used.

Morpheus is a full-featured P2P file-sharing system introduced by MusicCity (www.comusiccity.com) that tries to overcome some of the limitations of Napster. It includes a powerful search engine that can search for all types of media (including music, videos, documents, and reference files). The results are grouped together so the same file is displayed only once. Its SmartStream mechanism automatically resumes broken content streams by finding another source for the same content and monitoring the network until the whole content stream is downloaded. Morpheus increases the download speed of large files through the simultaneous transfer of content from multiple sources (FastStream mechanism). Its encryption mechanisms protect privacy and transmissions, and prevent unauthorized intrusions. Also, it allows content providers to deploy third-party digital rights management technology to protect the copyrights of their digital content that is distributed through the network.

Kazaa is another example of a P2P file sharing system that uses SuperNodes as local search hubs. These are powerful nodes on fast connections that are generated automatically. Peers connect to their local SuperNode to upload information about the files they share, and to perform searches in the network. Kazaa uses an intelligent download system to improve download speed and reliability. The system automatically finds and downloads files from the fastest connections, failed transfers are automatically resumed, and files are even downloaded from several sources simultaneously to speed up the download. When files are imported, the system automatically extracts meta-data from the contents of the files (such as ID3 tags for mp3 files). This makes for much faster and more accurate searches. Kazaa also uses a technique called MD5 hashing to make sure the contents of multi-sourced files are identical.

5.4 Collaboration

Collaborative P2P applications aim to allow application-level collaboration between users. The inherently ad-hoc nature of P2P technology makes it a good fit for user-level collaborative applications. These applications range from instant messaging and chat, to online games, to shared applications that can be used in business, educational, and home environments. Unfortunately, a number

of technical challenges remain to be solved before pure P2P collaborative implementations become viable.

Overview. Collaborative applications are generally event-based. Peers form a group and begin a given task. The group may include only two peers collaborating directly, or may be a larger group. When a change occurs at one peer (e.g., that peer initiates sending a new chat message), an event is generated and sent to the rest of the group. At the application layer, each peer's interface is updated accordingly.

Technical Challenges. There are a number of technical challenges that make implementation of this type of system difficult. Like other classes of P2P systems, *location* of other peers is a challenge for collaborative systems. Many systems, such as Magi, rely on centralized directories that list all peers who are online. To form a new group, peers consult the directory and select the peers they wish to involve. Other systems, like Microsoft's NetMeeting, can require that peers identify one another by IP address. This is much too restrictive, especially in environments where groups are large.

Fault tolerance is another challenge for collaborative systems. In shared applications, messages often must be delivered reliably to ensure that all peers have the same view of the information. In some cases, message ordering may be important. While many well-known group communication techniques address these challenges in a non-P2P environment, most P2P applications do not require such strict guarantees. The primary solution employed in P2P applications is to queue messages that have been sent and not delivered (i.e., because a given peer is down or offline). The messages can then be delivered to the offline peer when it comes back online.

Realtime constraints are perhaps the most challenging aspect of collaborative implementations. Users are the ultimate end points in a collaborative environment. As such, any delay can be immediately perceived by the user. Unfortunately, the bottleneck in this case is not the P2P technology, but the underlying network. While many collaborative applications may work well in a local-area systems, wide-area latencies limit P2P applications just as they limit client-server applications. Consider a gaming environment. The game DOOM is a so-called First Person Shooter (FPS) game in which multiple players can collaborate or compete in a virtual environment. DOOM uses a P2P structure in which each player's machine sends updates of the state of the environment (such as the player's movement) to each of the other machines. Only when all updates have been received does the game update the view. This was marginally viable in local-area, small-scale games, but did not

scale to wide-area games. Long latencies and uneven computing power at the various players machines made this lock-step architecture unusable. All FPS games since DOOM have used a more standard client-server architecture for communication.

5.5 Platforms

Operating systems are becoming decreasingly relevant as environments for applications. Middleware solutions, such as Java Virtual Machines, or Web browsers and servers are the dominant environment that is of interest to users as well as to developers of applications. In that regard, it is likely that future systems will increasingly depend on some other sort of platform that will be a common denominator for users and services connected to the Web or in an ad-hoc network. Examples of such environments include AOL and Yahoo, and .NET is striving toward a similar goal.

As described in Section 3.1, platforms, even more so than other P2P systems, have support for primary P2P components: naming, discovery, communication, security, and resource aggregation. They have an OS dependency even though it is minimal. Most P2P systems are either running on an open-source OS (Linux) or they are based on Windows.

There are a number of candidates competing for future P2P platform. .NET is the most ambitious one, going beyond P2P to encompass all service support on the client and server side. JXTA is another attempt, taking a bottom up and strong interoperability approach. Most other systems also have some level of platform support, such as Groove covering enterprise domain and Magi, covering handheld devices domain. Section 6.7 and Section 6.8 contain detailed descriptions of JXTA and .NET respectively.

6 CASE STUDIES

In this section, we compare eight case studies of P2P systems. We selected systems in four different categories, representing the spectrum of different P2P system categories, as well as public domain and proprietary systems (see Table 5).

6.1 Avaki

Avaki provides a single virtual computer view of a heterogeneous network of computing resources. It is a classic example of meta-computing applied to networks ranging from corporate compute and data grids to global application grids of Internet scale.

P2P system	Developer	Technology
Avaki	Avaki	distributed computing
SETI@home	public domain	
Groove	Groove Networks	collaboration
Magi	Endeavors Technologies	
FreeNet	public domain	content distribution
Gnutella	public domain	
JXTA	public domain	platform
.NET	Microsoft	

Table 5. Case Studies for Eight P2P systems

Figure 16: Avaki Architecture.

History. Avaki began as Legion [Grimshaw et al. 1994], a research project initiated by Andrew Grimshaw at the University of Virginia in 1993. The vision was to achieve a unified view of the computing resources scattered around the nation as a single nationwide virtual computer. Passing through various stages of development, and following the first showcase of the Legion technology in 1997 at the Super Computing conference, the research project emerged as a commercial venture called Applied MetaComputing, in 1998. In mid 2001, it was relaunched as Avaki Corporation, which currently focuses on providing enterprise-level distributed computing solutions.

Goals. The eventual goal is to knit together the nation's computing resources into a seamless parallel execution environment, allowing faster completion of applications. Current goals include robust security, performance management, and failure detection and recovery features to radically simplify grid administration. Avaki is marketed as a middleware platform for enterprise-level computing. They are also working with other developing standards in the field of distributed, pervasive, and P2P computing, such as JXTA, to make Avaki an interoperable platform.

Design. The core of Avaki is based on the object-oriented paradigm. Every entity in the system is an individually addressable object with a set of interfaces for interaction. This approach enables uniformity in system design through inheritance, and containment through encapsulation. Legion was an extension to Mentat [Grimshaw 1994], an object-oriented parallel processing system. The ability of Mentat to deliver high performance on platforms with very different communication characteristics was the key reason it was chosen to be extended to produce Legion. High performance, exploitation of heterogeneity, scalability, and masking of complexity were key design criteria in the development of Avaki.
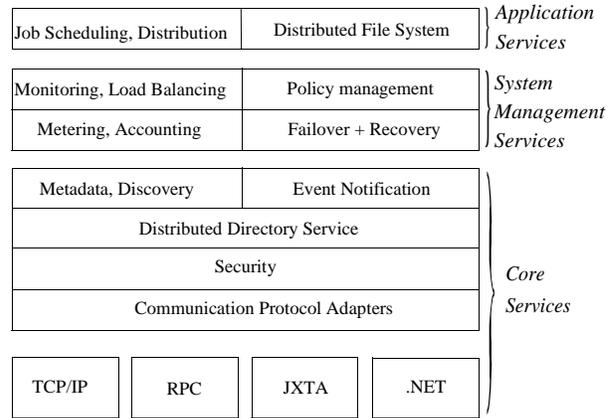
The middleware is structured as a layered virtual machine as shown in Figure 16. The stack has three main components.

- *Core Services*. These service provide some of the basic functionality required to map the system to a network of computing resources. The meta-data and the directory services enable users and applications to efficiently locate files or computing resources. The protocol adaptors enable interoperability with various networking standards including JXTA – a developing open standard from Sun.
- *System Management Services.* These services allow the Avaki system to monitor and control the efficiency of the meta-system. The policy management service allows administrators to manage access control for local resources.
- *Application Services.* Built over the basic services, these services enable the construction of applications such as collaborative and high-performance computing.

Scalability. Avaki is built as a layered virtual machine with scalable, distributed control. It allows for multiple administration domains, allowing system administrators complete control over local resources and their access policies. To make the file system scalable and to not impose a single file system across all machines, Legion takes a federated file system approach and builds a unified file system over existing ones.

Fault resilience. The scale of the Avaki network enables high redundancy of hardware components, but at the same time detecting and recovering from faults in this large system is a challenge. Legion decides to tradeoff

extensive fault tolerance support for higher performance. The scope of the support is limited to fail-stop faults of hardware components and chooses to stay away from software or Byzantine faults. However, in the event of internal failures during program execution, the user is notified and can decide on appropriate steps. When a host goes down, Avaki dynamically reconfigures the grid and automatically migrates failed jobs to different locations. Applications are allowed to specify the level of fault tolerance they require, the idea being that they should not pay for fault tolerance that is not required.

Implementation and performance. Because Avaki is designed as a layered virtual machine, the applications incur an overhead when compared to native executions. However, Avaki gains significantly by the parallel execution of applications, and trades off the overhead for significantly reduced software complexity. Applications submitted to the Legion system are executed in parallel across a pool of available computing resources. The main design goal is the reduction of completion time of the application. This obviously depends on how parallelizable the application is and the existing data dependencies. Applications that can tolerate latency and are coarse grained stand to gain from the parallel execution. Results from executing applications such as CHARMM and complib over Legion show an order of magnitude speedup over executions at a single installation.

Security. Avaki has built-in security, which eliminates the need for any other software-based security controls. Its security is lateral and decentralized allowing individual administration domains control over access to local resources. Authentication by the users occurs once during sign-in, and Avaki manages all further authentication required to run tasks over resources spread across multiple administration domains.

Lessons learned. Handling heterogeneity and automatic error detection and recovery are key to the success of systems of this scale. Systems such as Avaki also have to handle the possibility of a small percentage of the results being incorrect. One approach to handling this problem is to evaluate the same job at separate locations and verify the consistency of the results.

Business model. Avaki is marketing its product to enterprises as a complete data and compute grid. However, it faces tough competition from established HPC solution vendors such as Platform Computing. Avaki is currently being evaluated at various research labs.

Applications. Apart from enterprise solutions, Avaki can provide high-end computing power for problems in
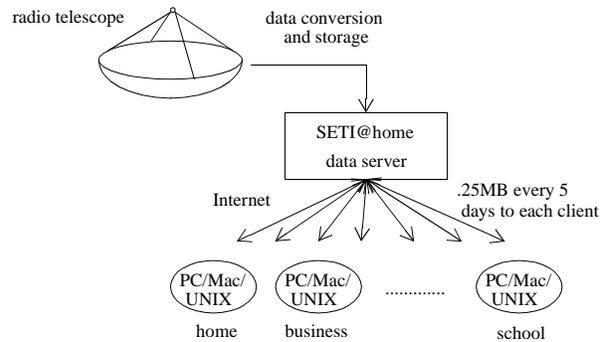

Figure 17: SETI@home Architecture. *Adapted from [Sullivan et al. 1997].*

the pure science area. Applications that involve parameter-space studies can benefit from the high-performance parallel execution characteristics of Avaki. Examples include biochemical simulations, *complib* – a protein and DNA sequence comparator and CHARMM – a p-space study of 3D structures.

6.2 SETI@home

SETI (Search for Extraterrestrial Intelligence) is a collection of research projects aimed at discovering alien civilizations. One of these projects, SETI@home, uses radio emissions received from space and collected by the giant Arecibo telescope, it analyzes them using the processing power of millions of unused Internet PCs.

History. SETI@home is a scientific research project aimed at building a huge virtual computer based on the aggregation of the computer power offered from internet-connected computers during their idle periods. The project has been widely accepted and raised an incredible raw processing power of several dozens of TFlops from more than three million internet-connected computers.

Goals. The goal is to process a search for extraterrestrial radio emissions from nearby developed intelligent populations based on data collected from the huge radio telescope at Arecibo. The ultimate objective of SETI@home, realized within the framework of the SETI project, is to track down and identify signals sent from intelligent civilizations situated outside our solar system.

Design. A major value of SETI@home is that the project old enough so that the tools have reached a very high level of quality. The project uses two major components: the database server and the client. They are provided by several different platforms, such as Windows, Linux, Solaris, and HP-UX (see Figure 17). The database has proven to be very scalable (more than three million users registered) and solid. Parts of the server-side code have

been made available for analysis. The client-side software is available as a screen saver module (although stand-alone operation is available) and is not linked to any third-party technology. This means that the client software has been developed specifically for each supported platform. Although the required development effort was costly, it made it easier to solve specific problems such as handling security (sandboxing of the execution, encryption of information, digital digests).

Fault resilience. Fault resilience is a major value of the client operation. Because one computing batch may require as much as ten hours to complete, it was necessary to ensure seamless recovery when the user logs in (which stops the SETI@home client), as well as when the machine is shut down. The SETI@home client resilience relies on a *checkpointing* mechanism, where the resumable dataset is saved on the hard disk every ten minutes. This means that each time a SETI@home computation is interrupted (by the user or a system failure), the computation will resume at the last saved dataset and proceed from there. This simple mechanism adds only a small processing overhead and little complexity, and results in a program that is very reliable.

Scalability. The SETI@home project strongly relies on its server to distribute jobs to each participating peer and to collect results after processing is done. The horizontal scalability (number of users) of the system is excellent, with millions of enthusiastic users already participating. A *vertical* scalability bottleneck of the architecture may be that a single server is responsible for coordinating all of its operations. However, the huge number of users shows that the system can handle this load.

Lessons learned. The major value of the SETI@home project – apart the scientific results – was to prove that the technology can be applied to real situations. We can envision how it could be used to better tap the processing power from unused computers and identify the peak periods of used computers. SETI@home developers were hoping for approximately 100,000 participants, but they surpassed this number within a week. So far there have been more than 3,000,000 contributors. The last lesson is related to security. Most of the problems were not caused by malicious users, but as a consequence of competitiveness. People raised the frequency of their CPUs, making them more exposed to failures [Wortheimer 2002].

Business model. However, it is still not clear how vendors will package this approach into products and services, or what business model they should follow.
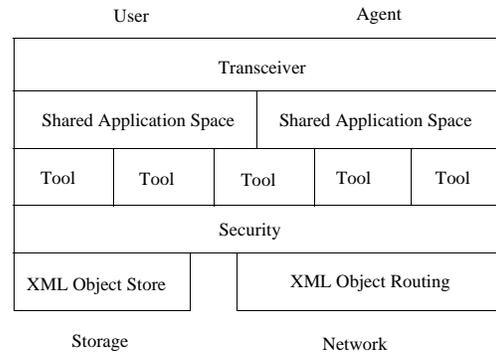

Figure 18: Groove Layers. *Adapted from Groove Product Backgrounder White Paper, http://www.groove.net/pdf/back-grounder-product.pdf.*

6.3 Groove

Groove is a collaborative P2P system. It is mainly targeted to Internet and intranet users, although it can also be used on mobile devices, such as PDAs, mobile phones, and tablets. It is intended to enable communication, content sharing, and tools for joint activities [Leigh and Benyola 2001].

History. Groove was founded in 1997 by Ray Ozzie, the developer of Lotus Notes. The first version of the product was released in the beginning of 2001.

Goals. The Groove's main goal, which was P2P by design, was to allow users to communicate directly with other users without relying on a server [Suthar and Ozzie 2000]. Other important goals include *security and privacy,* and *flexibility*.

Groove users are authenticated, data is secured both on disk and on the wire, and data confidentiality and integrity are enabled by secure key technology. User data residing on the server is opaque and private data is not shared with third parties. Finally, Groove components are signed.

Flexibility in Groove is related to both network and development. Groove supports the Internet, intranets, and private networks, and allows users to transparently handle disconnection. Groove also supports a number of reusable components.

Applications. Groove is intended for communication, collaboration, and content sharing.

- *communication*: voice over Internet, instant messaging, chat, threaded discussions
- *content sharing*: shared files, images, and contacts
- *collaboration*: group calendaring, group drawing and editing, and co-Web browsing
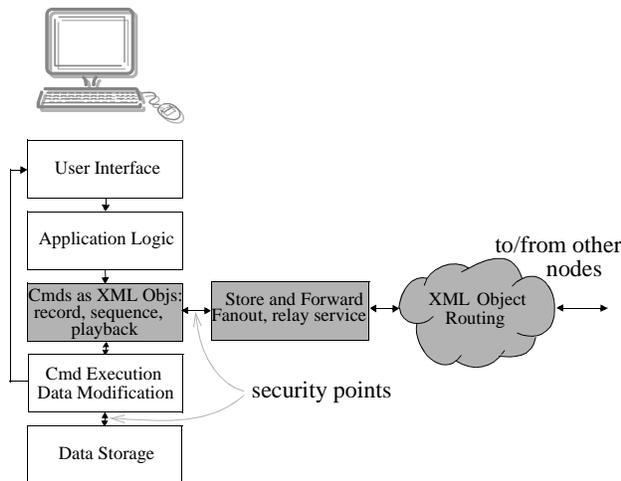
28

Figure 19: Groove Application Structure. *Adapted from http://www.groove.net/developers/presentations/architecture.exe.*

Groove presents itself to users or to agents representing users in the form of shared spaces (See Figure 18). Shared spaces offer users benefits, such as *spontaneity* – users act without an administrator, *security* – shared spaces act as virtual private networks, *context* – spaces provide and maintain the context for users, *synchronization* – all spaces of the same user are synchronized across all users devices, and *granularity* – parts of documents can be exchanged [Udell, et all 2000].

Design. The Groove layer is inserted between the application logic and command execution layers (see Figure 19). This enables the command requests or data to be stored locally or forwarded to the network, as appropriate. Commands are transformed into XML objects, recorded, and sequenced, which enables them to be played back and relayed in case of disconnection. Security layers are executed before storing objects to disk or sending them out to the network, preventing unauthorized access. Groove supports peer-based authentication and end-to-end encryption.

Groove supports system and centralized services [Groove Networks 2001]. *System services* include:

- *security*, such as automatic handling of public/private key management, encryption, and authentication
- *storage* based on the XML object store, enabling disconnection from the network
- *synchronization* of multiple local copies of the user spaces
- *peer connection* to support transparent administration, such as IP addresses, bandwidth selection, and firewall translators

Centralized services, which leverage centralized management of resources, include:

- *licence management* which enables viewing all the devices using certain software as well as managing the use of this software
- *component management,* to enable managing Groove components on the fly without requiring user interaction for upgrades
- *relays and transparent cross-firewall interaction* to optimize the communication bandwidth and offline work, while still enabling direct P2P communication
- *usage reporting* to track how shared spaces are used, enabling administrators to optimize resource distribution while respecting user privacy

Fault Resilience. The only form of fault resilience support is through relay activities, where messages sent to failed nodes are queued until the nodes are brought up. The same applies to transient failures in the network connectivity.

Business model. Groove Networks plans to license its infrastructure platform to corporations and third-party integrators [Rice and Mahon 2001]. The competitors are primarily other collaborative P2P systems, such as Endeavors Technologies, but also Microsoft, given the nature of the systems platform that both Groove and Microsoft advocate. Some of the enablers for adoption of Groove include elimination of the network administration costs, minimization of dependences on server infrastructure, and availability.

6.4 Magi

Magi is a P2P infrastructure platform for building secure, cross-platform, collaborative applications. It employes Web-based standards such as HTTP, WebDAV, and XML to enable communication among applications within an enterprise network or over the Internet. Magi Enterprise, their end product, builds over the infrastructure to link office and project teams so they can share files, do instant messaging, and chat.

History. Magi evolved from a research project headed by Greg Bolcer at the University of California, Irvine. His team was funded by grants from DARPA, and at the time, it was believed to be the largest, non-Sun Java project in the country. Endeavors Technology was founded in 1998, and was later bought over by Tadpole Technology. The first version of Magi, which is their P2P infrastructure, was released in late 2000. This was followed by their enterprise edition in 2001.
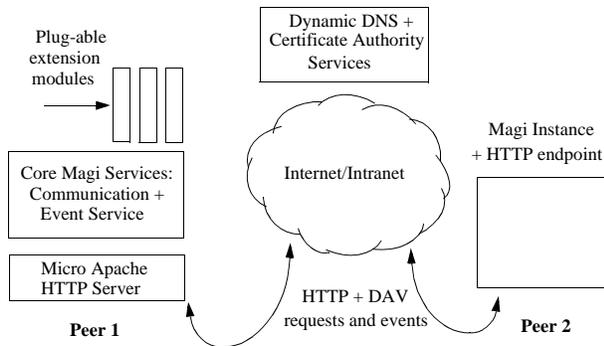
Figure 20: Magi Architecture.

Goals. The goal of Magi is to enable information sharing and messaging on any device using popular Web-based standards. The architectural framework runs on various platforms including PDAs and Pocket PCs. It aims to make development of XML and Java-based distributed applications easier by providing event notification to any device running an instance of Magi.

Design. The key components of the Magi framework include a micro-Apache HTTP server which provides a link to every instance of Magi, a set of core services, and a generic extensible interface. The modular architecture of an Apache server allows easy reconfiguration and the addition of interfaces only if they are required. The core services include:

- *Communication Service:* messaging interface between Magi services and the HTTP server

- *Event Service:* local modules register with the event service to receive specific event notifications; also enables communication among local modules

- *Buddy Manager:* maintains a list of active buddies and their most recent communication endpoints

- *Access Controller:* allows restricted access to resources on a Magi peer

Plug-able service modules such as MagiChat, Magi-DAV, and instant messaging, could be built over the core infrastructure as shown in Figure 20.

Each instance of Magi acts as both a client as well as a server and hence can participate in P2P communication with other Magi servers. Magi supports the following Web protocols: HTTP1.1; WebDAV for remote authoring of documents and SWAP/Wf-XML for remote process monitoring and control. The modules supporting these protocols can be loaded when required to achieve interoperability with applications compliant with these protocols. These features reinforce Magi's primary design goal to achieve interoperability with existing Web access standards.

Scalability. Magi relies on a centralized Dynamic DNS as a directory for IP addresses of Magi instances. Because Magi instances can have dynamic IPs, when a Magi instance comes up, it reports its IP to the DDNS. This enables other Magi instances to find the current network endpoints of their buddies. This centralized repository might be a scalability issue with the growing number of users on the Magi network. More importantly, it is a single point of failure, which could cause non-availability of aliasing information. For user authentication, Magi provides a centralized Certificate Authority (CA), which is another scalability concern and single point of failure.

Fault Resilience. Queuing messages that cannot be delivered to buddies who are currently offline enables guaranteed delivery when the destination buddy eventually comes online. Buddy crashes are not detected, as there is no mechanism to check the liveliness of each buddy. However, the core Magi interface can be extended as desired to support any kind of fault resilience mechanism among the buddies. For its centralized services such as dynamic DNS and CA, Magi currently does not provide any fallback schemes in the event of failure.

Implementation. The entire infrastructure is in Java and the Web interface is through servlets. This makes Magi as fast as the underlying virtual machine. Because each Magi instance is both able to receive and send messages over HTTP, a minimal Web server must be running. This may not be the best solution for a Pocket PC or a handheld with limited memory resources. Magi uses Tomcat, which supports modular loading of essential components, coupled with independent service-oriented design of the infrastructure to target constrained and embedded devices. Magi's services are accessible through a Web-based interface and is the access mechanism in their implementations for the Compaq iPaq and the HP Jornada pocket PCs.

Lessons learned. While most other P2P infrastructures implement proprietary interfaces and communication standards, Magi emphasizes the use of existing popular standards. This makes Magi a highly interoperable platform, and its Web-based design makes its deployment on a range of devices easy.

Applications. Magi was designed primarily with the paper-based workflow scenario in mind and is targeted at any type of collaborative environment. It supports file sharing and collaborative tools such as chat and messag-

ing. As a platform, it can be used to embed collaborative processes into enterprise-wide applications and B2B products. An SDK has also been released to ease the integration of Magi into applications.

6.5 FreeNet

Freenet is a P2P file-sharing system based on an initial design by Ian Clarke [Clark 1999, Clark et al. 2001]. The primary mission of Freenet is to make use of the system anonymous. That is, upon entering the system, a user should be able to make requests for files without anyone being able to determine who is making these requests. Likewise, if a user stores a file in the system, it should be impossible to determine who placed the file into the system. Finally, the operator of a freenet node should have no knowledge of what data is stored on the local disk. These forms of anonymity make it possible to provide storage and use the system without concern of being tracked down or potentially held liable.

History. The Freenet system was conceptualized by Ian Clarke in 1999 while at the University of Edinburgh. In his introduction to Freenet, he cites the following quotation:

"I worry about my child and the Internet all the time even though she's too young to have logged on yet. Here's what I worry about. I worry that 10 or 15 years from now, she will come to me and say "Daddy, where were you when they took freedom of the press away from the Internet" - Mike Godwin

Public development of the Open Source Freenet reference implementation began in early 2000.

Goals. The principle goal of Freenet is to provide an anonymous method for storing and retrieving information. Freenet permits no compromise on these goals. However, within these bounds, Freenet also strives to be as reliable, easy to use, and responsive as possible.

Design. One of the key design points of the Freenet system is to remain completely decentralized. Therefore, Freenet represents the purest form of P2P system. Freenet's basic unit of storage is a file. Every node in the Freenet network maintains a set of files locally up to the maximum disk space allocated by the node operator. When all disk space is consumed, files are replaced in accordance with a least recently used (LRU) replacement strategy.

Each file in the Freenet system is identified by a key. These are typically generated using the hash SHA-1 [SHA-1 1997] function. A variety of mechanisms are used to generate the desired hashes, but typically a user starts by providing a short text description of the file. This description is then hashed to generate a key pair. The public key becomes the file identifier. The private key is used to sign the file to provide some form of file integrity check. Other schemes for generating keys can be used as well permitting users to create hierarchical file structures or to generate disjoint name spaces. The file's key is then made available to users of the system by out-of-band mechanisms such as a Web site. Because the key can be computed from the description of the file, it is common to publish only the description and not the actual key.

The only operations in the Freenet system are inserting and searching for files. In either case, it is essential to find the proper location for the file. In general, Freenet nodes form a network in which they pass and forward messages to find the location of an existing file or the proper location to store a new file. The keys are used to assist in the routing of these messages. Freenet attempts to cluster files with similar keys on a single node. By clustering, Freenet is able to optimize searches by creating routing tables. When a file is successfully located by a search, the file's key is inserted into a local routing table. When another search message is received, it is first forwarded to the peer node with the most similar key in the routing table. When a search is received by a node that contains the desired file, it returns the entire file as a successful result. This is done recursively until the file is returned to the initial requester. As a side effect, the file becomes replicated at each node in the search path. In this way, popular files become highly replicated.

Inserts operate much the same as searches. First, a search operation is performed on the file's key. If a file with that key is found, it is returned to the inserter as a form of key collision indication. When no existing file is found, the file is propagated forward along the search path. This accomplishes both the replication of the file that occurs during a search as well as preserving the integrity of the routing tables by placing the new file in a cluster of files with similar keys.

New nodes announce their presence in the network by performing a search operation. This search basically accomplishes the function of announcing the new node to other existing nodes. Prior to sending the message, the node must first discover at least one other node in the network to which it can connect. Freenet explicitly does not help in this problem because doing so typically leads to centralized solutions. User's are required to bootstrap

themselves into the network using out-of-band means for finding at least one peer.

Scalability. The scalability of Freenet has been studied by its authors using extensive simulation studies [Clarke et. al. 2001]. Their studies support the hypothetical notion that route lengths grow logarithmically with the number of users. In their experiments, they start with a network of 1000 nodes in a ring topology. As the experiment runs, random keys are inserted and removed from the network, and the path length to find a file reduces from approximately 500 to less than ten. This is consistent with the logarithmic curve expected.

Implementation and performance. Freenet is available in an Open Source reference implementation. The protocol is also well defined allowing others to create their own implementations. One side-effect of the focus on anonymity in Freenet is the difficulty in observing its behavior. Obscured identities and probabilistic choices make measurements difficult. For these reasons, no real world performance studies seem to be available. Only the simulation results described above can be used to evaluate the system.

Lessons learned. The key lesson of Freenet is both the importance and difficulty of maintaining anonymity. Anonymity opens the door to freer discourse on the network because users need not be concerned with the ramifications of their actions. Anonymity also runs contrary to many intellectual property notions such as copyright. The Freenet team argues that preserving freedom and eliminating censorship are more important than intellectual property concerns.

Business model. Freenet operates as a non-profit Open Source project. There are as yet no commercial ventures building on top of Freenet. The Freenet project does solicit donations to help fund continued development, but does not appear to be dependent on these donations.

Applications. Freenet's only real application is as an information storage and retrieval system. The heavy use of cryptography as well as anonymization of requests may lead to other related uses. For example, it may be possible to use Freenet as a form of distributed backup system with some guarantees that data can only be retrieved by the owner of the data. However, Freenet only provides probabilistic guarantees about the persistence of data, so this likely would not make for a high-confidence back-up solution.

6.6 Gnutella

Gnutella is a file sharing protocol. Applications that implement the Gnutella protocol allow users to search for and download files from other users connected to the Internet.

History. Gnutella file sharing technology [Gnutella, 2001] was introduced in March of 2000 by two employees of AOL's Nullsoft division. Touted as an open-source program with functionality similar to that of Napster [Napster, 2001], the Gnutella servant program was taken offline the following day because of a possible threat to Warner Music and EMI. AOL was rumored to be in the midst of merger talks with the record companies at that time. However, the open-source program remained online long enough for eager hackers to discover the Gnutella protocol and produce a series of clones to communicate using the Gnutella communication protocol. Soon after, versions of the original Gnutella servant were communicating with Gnutella clones to search and trade files over the Gnutella Network.

Goals. The goal of Gnutella is to provide a purely distributed file sharing solution. Users can run software that implements the Gnutella protocol to share files and search for new files. The decentralized nature of Gnutella provides a level of anonymity for users, but also introduces a degree of uncertainty.

Design. Gnutella is not a system or a piece of software. Gnutella is the communication protocol used to search for and share files among users. A user must first know the IP address of another Gnutella node in the network. This can be discovered by going to a well-known Web site where a number of Gnutella users are posted. When a user wishes to find a file, the user issues a query for the file to the Gnutella users about which it knows. Those users may or may not respond with results, and will forward the query request to any other Gnutella nodes they know about. A query contains a Time-To-Live (TTL) field and will be forwarded until the TTL has been reached.

Scalability. While the Gnutella model has managed to succeed thus far, in theory it does not scale well. The number of queries and the number of potential responses increases exponentially with each hop. For example, if each node is connected to only two others and the TTL of a query is 7 (the default for most Gnutella queries), the number of queries sent will be 128 and the number of responses may be substantially more depending on the popularity of the item.

Fault resilience. The Gnutella protocol itself does not provide a fault tolerance mechanism. The hope is that

enough nodes will be connected to the network at a given time such that a query will propagate far enough to find a result. However, the distributed nature of the protocol does not guarantee this behavior. In fact, studies have shown [Adar and Huberman, 2000, Saroiu et al. 2002] that only a small fraction of Gnutella users actually remain online long enough to respond to queries from other users. More robust software has begun to alleviate this problem, however no guaranteed solution exists. The only real solution at this point is to rely on users to retry when their queries or downloads fail.

Implementation and performance. Since the first Gnutella servant was posted by Nullsoft, a number of companies have implemented clone software and made efforts to overcome many of the limitations not addressed by the protocol. Among the most popular are Limewire [Limewire, 2001], ToadNode [ToadNode, 2001], and BearShare [BearShare, 2001]. Kotzen claims that while in late 2000 only 10% of download attempts were successful, by March of 2001 the number grew to over 25% [Kotzen, 2001]. In addition, the number of Gnutella users has increased. However, Kotzen reports that a maximum of 11,000 users has been seen on the network at any one time. While this represents an increase over previous findings [Clip 2 Marketing, 2000], it does not provide any proof of the scalability or performance of the Gnutella network for the targeted thousands or millions of nodes.

Business model. Gnutella is not a company or a piece of software, but rather an open protocol. As such, there is no Gnutella business model. Many of the companies developing Gnutella-compatible software are open-source or offer free download. Additionally, some companies are building software for the enterprise that offer content-specific content sharing.

Lessons learned. The Gnutella model has raised a number of questions in the research and development community. Well-known solutions to traditional computer systems problems cannot be applied in a straightforward manner to Gnutella's decentralized network. Researchers are busily trying to apply distributed computing and database problems to ad-hoc, unorganized, decentralized networks of nodes. The benefit of a Gnutella-style network is the use of an unbounded set of diverse resources. However, solutions often trade reliability for the potential of discovering an otherwise unavailable resource.

Applications. The primary application for this technology has been the sharing of music files. While this continues to be the main motivating application for companies developing Gnutella-compatible software, Gnutella is more a paradigm than a given application. A number of

companies are developing visions of using the Gnutella protocol for enterprise software including project-management-style applications as well as academic software for classroom-based communication.

6.7 JXTA

The vision of the JXTA project is to provide an open, innovative collaboration platform that supports a wide range of distributed computing applications and enables them to run on any device with a digital heartbeat. JXTA provides core functionality in multiple layers, including basic mechanisms and concepts, higher level services that expand the capabilities of the core, and a wide range of applications that demonstrate the broad applicability of the platform.

History. The JXTA project was unveiled by Sun on April 25, 2001 [JXTA 2001] and was intended to be a platform on which to develop a wide range of distributed computing applications. Despite its recent introduction, JXTA has been quite popular. Statistics show that on the week of November 3rd, 2001, JXTA had 122 posts and was used by 6,809 users.

Goals. The goal of JXTA is to provide a "general-purpose" network programming and computing infrastructure. Its goals are:

- *Interoperability*: by enabling inter-connected peers to easily locate each other, participate in community-based activities and offer services to each other seamlessly across different P2P systems and different communities

- *Platform independence*: JXTA is designed to be independent from programming languages (such as C or Java), system platforms (such as Microsoft Windows and UNIX operating systems), and networking platforms (such as TCP/IP or Bluetooth)

- *Ubiquity:* JXTA is designed to be implementable on every device with a digital heartbeat, including appliances, desktop computers, and storage systems

Design. It is important to note that the JXTA project is approaching the P2P space from the lower level as they are proposing an entirely new infrastructure with no direct relation to other, existing P2P systems (e.g., Gnutella and Napster). For example, they have built their own distributed search service, called JXTA search.

Peer groups are the core of JXTA's infrastructure. A peer group is essentially a partitioning of the world of peers for communication, security, performance, "logical locality" and other reasons. A single participant can be in multiple groups at one time. JXTA provides core proto-

| JXTA
Applications | | Shell
InstantP2P |
| JXTA
Services | Indexing
Searching | Content
Manager |

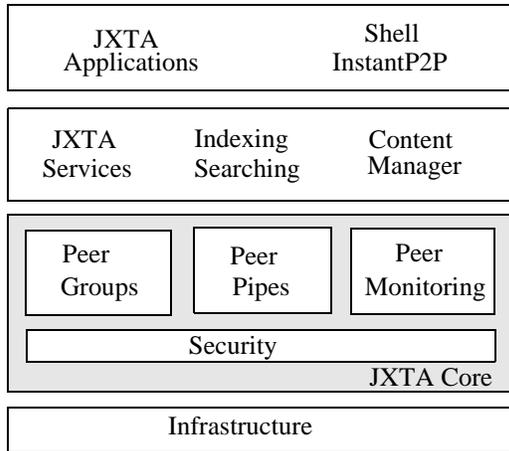| Peer
Groups | Peer
Pipes | Peer
Monitoring |
| Security | | |
| | | JXTA Core |

| Infrastructure |

Figure 21: JXTA Architecture.

cols for peer discovery, peer group memberships, and peer monitoring. JXTA uses asynchronous uni-directional communication channels, called *pipes,* for sending and receiving messages. All data interchange in JXTA is in the form of XML formatted documents.

Services. The JXTA project has defined core and optional services that run on the JXTA platform. Examples of core services include authentication, discovery, and management. To address the need for security, they have implemented a cryptography toolkit that enables message privacy, ensures authentication and integrity, and permits transactions between peers that cannot be repudiated. However, it is not clear at this point how they deal with firewalls; for example, how a peer outside a firewall can discover peers inside the firewall or how two peers inside two firewalls can communicate. Another core service is the Content Manager Service (CMS) that allows JXTA applications to share, retrieve, and transfer content within a peer group. However, in the current version of the CMS, there is no support for automatic propagation of content requests; each peer sends content requests directly to the other peers. Examples of optional services are naming, routing, indexing, and searching. For example, the JXTA distributed indexing service implements a general-purpose, fully distributed index service by using a distributed inverted index to map keywords to a list of postings.

Applications. The JXTA Shell was the first application developed that gives a command-line interface to the core JXTA services, such as peer and group discovery and messaging. Since then, the JXTA project has developed more applications that allow interactive access to the JXTA platform. For example, the InstantP2P application enhances the JXTA project with chat capabilities. It includes functionality that enables users to chat one on one, chat with a group, and share files. Another applica-

tion is JuxtaProse, which is a discussion application that allows creating, viewing, and replying to HTML documents over JXTA's content management system. Other applications that the JXTA community is currently designing include event notification and P2P email.

Scalability. The JXTA project has a layered architecture, which consists of three building blocks: JXTA core (for peer group management), JXTA services (such as searching and indexing), and JXTA applications (e.g., JXTA Shell). This layered architecture enables JXTA to easily incorporate new protocols and services to support a growing number of users. However, it will certainly face some scalability problems that have been currently left un-attended. One of these problems is global naming. JXTA attaches a unique ID (generated by the users) with each peer in the group but does not solve the global naming problem because it does not guarantee uniqueness across the entire community of millions of peers. Even if a peer has a unique ID within its group there is no guarantee that its name is unique across multiple groups.

Implementation and performance. The first version of the JXTA core has currently been released, running in three different platforms: the Solaris Operating System, Linux, and Microsoft Windows. A prototype of the JXTA Shell that gives a command-line interface to the core services is also available. In addition to that, there have been many independent efforts that build services and applications on top of JXTA.

Business model. JXTA has been released as open-source code and has already attracted many developers to build applications (e.g., event notification, file sharing, P2P email) on top of it. Sun is also working with other P2P companies that are committed to work with the JXTA open-source model and contribute their P2P technology to it. JXTA requires some familiarity with the UNIX operating system, which is not required for typical Gnutella users.

6.8 .NET My Services

.NET My Services and .NET in a more global form do not represent a P2P system in its traditional form, as do other P2P systems described in this paper. However, .NET My Services encompass a lot of P2P architecture and as such they are presented here. .NET also introduces a new language called C#, development tools, a framework, a foundation for Web services, and the underlying operating systems. This paper addresses only the .NET My Services part of it (shaded part of Figure 22).
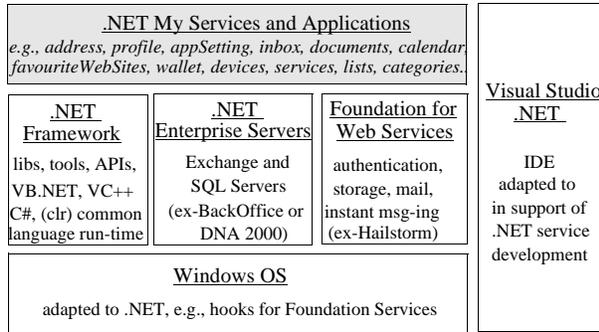
.NET My Services and Applications
e.g., address, profile, appSetting, inbox, documents, calendar favouriteWebSites, wallet, devices, services, lists, categories.

.NET Framework	.NET Enterprise Servers	Foundation for Web Services
libs, tools, APIs, VB.NET, VC++ C#, (clr) common language run-time	Exchange and SQL Servers (ex-BackOffice or DNA 2000)	authentication, storage, mail, instant msg-ing (ex-Hailstorm)

Windows OS
adapted to .NET, e.g., hooks for Foundation Services

Visual Studio .NET

IDE adapted to in support of .NET service development

Figure 22: .NET Architecture.

History. Microsoft officially announced .NET for the first time in June 2000. Shortly thereafter, certain components were available, such as operating systems, parts of the .NET framework and enterprise servers, and passport. During 2001 and 2002, subsequent versions of operating systems, frameworks, and some key services have been introduced.

Goals. The goals of .NET My Services and .NET in the broader context are to enable users to access Web services on the Internet from available devices, including desktops, laptops, and handhelds, using existing standards, such as XML [Bray et al. 2000], UDDI [Ariba et al. 2000], SOAP [Box et al. 2000], WSDL (Christensen et al. 2001]. The goal of .NET is to shift focus from applications, platforms and, devices to user-focused data. Therefore, security and privacy are one of the major goals of .NET.

Design. The design of .NET is focused around de-componentization and decentralization of distributed services. For example, a user will locate the desired service through the UDDI registry, then she will access service through the Web Services provider location, and she will be authenticated through a separate passport service (see Figure 23). The .NET Services programming model is built around Web standards and standard protocols, such
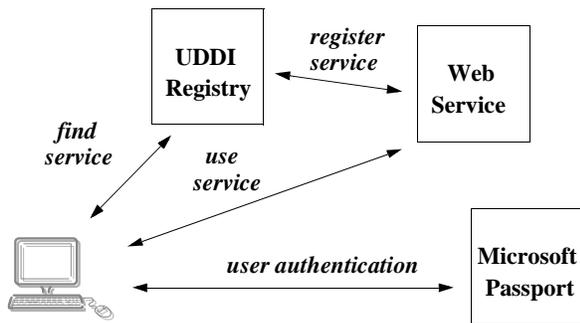

Figure 23: User Interaction with .NET Components.

as SOAP, UDDI, and WSDL. The .NET System programming model is built around a new common language run-time, which is intended to complement the Java run-time. Passport (see Security and Privacy below) is an authentication method for .NET Services. Other Web services can communicate with .NET My Services without disclosing a user's identity and personal information.

Security and Privacy. Security and privacy is one of the main goals but also one of the main challenges for .NET My Services. The support is provided through an online service (Passport) that enables authentication of users. This is used for accessing Web pages (see Figure 24) or for services in general. Once authenticated, the user can roam across passport-participating Web sites. Web sites in turn are required to implement the "single sign-in" service. Passport was released in July 1999 [Microsoft 2001a] and it had over 160 million accounts by July 2001.

Users of passport are required to create a passport account that contains the following information.

- *Passport unique identifier (puid).* Created automatically as a part of account creation.
- *User profile.* Contains an email address or a phone number (required), a first/last name, and demographic information (postal code, region, country, etc.).
- *Credential.* Consists of a standard credential (email/phone and a password) and a four-digit key for strong credential sign-in required by some Web sites.
- *Wallet* (optional). Contains credit card numbers and a postal address for delivery of purchased goods.

Of the passport information, only the puid is shared with participating Web sites, and wallet information is shared in case of purchases. This protects the privacy of users, with passport playing the intermediary in user authentication. Passport offers the Kids Passport service, where parents can restrict the information visiting Web sites can collect about kids. Finally, passport has been optimized for phone access by means of adaptation to minimal screen requirements and different protocol support (e.g., WML for WAP browsers or cHTML for i-mode). Passport will be integrated with Windows XP by storing them in the local credentials manager.

Scalability. It is only fair to compare the scalability of .NET implementations with its earlier pre-.NET equivalents. Operating systems will be enhanced with hooks for the benefit of Web Service Foundations, and the OS server editions will be required to scale even more given the increased number of clients with the introduction of mobile users that may have unpredictable access pat-

P2P System	System Feature					
	System Category	Alternative Solution	Platform	Languages and Tools	Distinctive Features	Target Networks
Avaki	distributed computing	single installation HPC and supercomputers	Linux, Solaris, MS Windows	OO, paral. lang. obj wrap., legacy app x-comp. (Fortran, Ada, C)	dist. admin ctrl, heterogeneity, secure access, high paral. exec	Internet, intranet
SETI@home		distributed objects	all common OS supported	closed source	large scale	Internet
Groove	collabora-tion	Web-based collaboration	MS Windows	JavaScript, VB, Perl, SOAP, C++, XML	exec. playback, self-updating multi-identities	Internet, intranet
Magi		dist. file sys. centralized chat & messaging	Windows and Mac	Java, XML, HTTP, WebDAV	HTTP based, platform indep.	Internet, ad-hoc (mobile)
FreeNet	content distribution	anonymity: none. others cen-tralized & single point of trust	any with Java	Java implementation and APIs	Preservation of anonymity	Internet
Gnutella		central servers	Windows, Linux	Java, C	protocol	Internet
JXTA	platform	client-server	Solaris, Linux, MS Windows	Java , C, Perl	open-source effort	Internet
.NET / My Services		Web-based	Windows	C#, VC++, JScript, VBScript, VB	widespread MS apps base	Internet and mobile

Table 6. Summary of Case Studies.

terns. The ultimate benefit of the .NET approach will be in the further decomponentizing of the solutions en-abling finer granularity of concurrency between compo-nents and hence better component utilization. At the same time, the side effect is improved reliability and availability.

Fault resilience. Similar to scalability, further decompo-nentization of .NET enables fewer single points of fail-ure and more opportunities for failover among peer components.

Implementation and performance. It is fairly early to talk about implementation details and performance. So far, most of .NET is envisioned to execute on Windows platforms. It is still uncertain whether there will be other platforms where .NET can run. Given the large base of developers it is likely that performance will be addressed as the deployment progresses.

Business model. .NET is a significant move in Mi-crosoft's strategy. It represents a move from a license-based to a subscription-based model. .NET also has sig-nificant impact on how other businesses will function in general. It is one of the first times that Microsoft embrac-es the Internet model in its entirety. Microsoft plans to run .NET Services as a business, similarly to AOL or Ya-hoo. It will charge consumers and service providers for using .NET My Services. In the longer term, Microsoft may provide enterprise-level .NET My Services. One

important aspect of the business model is that Microsoft embraced open-source standards, such as XML, SOAP, WSDL, and UDDI, as well as its submission of the C# and Common Language Infrastructure to standards body European Computer Manufacturers Association (EC-MA).

Applications. The major advantage of .NET My Servic-es compared to other solutions is the wealth of applica-tions. Microsoft will port all of the Microsoft Office applications to .NET, including the development tools.

Lessons learned. .NET and .NET Services are fairly new to be able to derive any substantial lessons learned. Nevertheless, at least two lessons can already be derived. First, the passport model raises a lot of concerns about security problems, but it has also attracted a huge number of users so far. Second, it is the applications that attract users; .NET seems to follow the same model as with MS DOS, where it is really the applications, such as Excel and Word that are of interest to users, not the infrastruc-ture itself.

6.9 Summary

This Section summarizes the case studies. Table 6 sum-marizes the general characteristics of all of the case stud-ies. The entries are self explanatory. P2P systems are deployed on a variety of systems and support different languages and standard protocols. They are largely tar-geted for the Internet and some are targeted for mobile settings.

Table 7 summarizes the case studies from the perspec-tive of the characteristics described in Section 4. There are various decentralization models, including pure (FreeNet, Gnutella, JXTA), hybrid (Groove, Magi), mas-ter-slave (SETI@home), and mixed (.NET). Few sys-
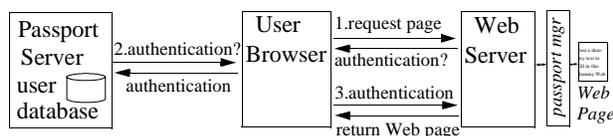

Figure 24: .NET Passport.

P2P System	System Feature										
	Decentralization	Scalability	Anonymity	Self-Organization	Cost of Ownership	Ad-hoc	Performance	Security	Transparency	Fault Resilience	Interoperability
Avaki	distributed, no central mgr	scale to 1000s 2.5-3k tested	N/A	restructures around failure	low	join/leave compute resources	speedups	encryp, authentication, adm. domains	location; HW/ SW heterog.	chckpt/restart reliable msg	interoperates with Sun Grid
SETI@home	master-slave	millions	medium	low	very low	join/leave compute resources	huge speedups	proprietary	high	timed chckpt	IP?
Groove	hybrid P2P	N/A	poor	high	low	join/leave of collaborators	medium	shared-spaces authn/authr, encryp	high	queued messages	IP-based
Magi	hybrid P2P	around a 100 (corporate NW)	N/A	N/A	low	join/leave of buddies	N/A	certificate authority	communicate offline buddies	queued messages	JXTA, WebDAV
FreeNet	pure P2P	theoret. scales ~ log(size_network)	high	high	low	join/leave of peers	medium	f anonymity & preventing DoS	high	No 1 point of failure, replic.	low
Gnutella	pure P2P	thousands	low	high	low	join/leave of peers	low	not addressed	medium	resume download	IP?
JXTA	pure P2P	also addresses embedded systems	N/A	N/A	low	join/leave of peers	N/A	crypto algor. distr. trust model	low	low	low
.NET / My Services	mixed	world-scale	N/A	medium	low	join/leave of peers	high	passport-based	high	replication	SOAP, XML, UDDI, WSDL

Table 7. Comparison of Characteristics of Case Studies.

tems support a high degree of anonymity (FreeNet). A number of them support some sort of self-organization and most of them improve cost of ownership. Ad-hoc nature is predominantly supported through the ability of the peers to join and leave. Performance is improved primarily in the distributed computing type of P2P systems (Avaki, SETI@home). Standard security mechanisms are supported, such as encryption, authentication, and authorization, and some of them protect against denial of service attacks (Freenet).

Transparency is primarily supported for the benefit of disconnection and to hide communication/computation heterogeneity (Groove, Avaki). Fault resilience is supported through checkpoint-restart mechanisms (Avaki, SETI@home), queued messages (Groove, Magi), and in general by the redundancy of components and the elimination of a single point of failure. Most systems interoperate by relying on standard communication stacks (IP

and on Web standards (SOAP, XML, UDDI, WSDL, and WebDAV). There are a only a couple of exceptions of interoperability with other P2P systems, such as Avaki with Sun's Grid, and Magi with JXTA.

Table 8 summarizes the business aspect of the case studies. The surveyed case studies belong both to proprietary (Groove, .NET) and open source systems (FreeNet, Gnutella, JXTA), or they are offered as either (Avaki, Magi). They support applications in their respective domains, such as distributed computing, communication/collaboration, and file sharing. P2P systems classified as platforms (.NET My Services, JXT, and to a certain extent Groove and Magi) support a wider base of applications.

The case studies we selected have a closed base of customers with the exception of .NET My Services, which relies on its earlier customer base. Also, they are competitors within each class of P2P systems. These systems have different funding model: startups (Magi, Freenet,

P2PSystem	System Feature					
	Revenue Model	Supported Applications	Known Customers	Competitors	Funding	Business Model
Avaki	product and open-source	computation grid, shared secure data access	none. evaluated at several life sciences labs	Platform Computing Globus	startup	N/A
SETI@home	Academic	Closed	Academic	cancer@home fight_aids@home	government	sell more computers ads on screen savers
Groove	product	purchasing, inventory, auctions, etc.	N/A	Magi	IPO	collaborative tool of choice next Lotus Notes
Magi	product & open-source	Shared file access, messaging, chat	Global eTech, InterPro Global Part., Mongoose T., Mediasoft	Groove	startup	N/A
FreeNet	open-source	file sharing	public	N/A	startup	N/A
Gnutella	open-source	file sharing	public	N/A	public domain	algorithm of choice for P2P
JXTA	open-source proprietary extensions	file sharing, messaging, event notific., email	Many P2P systems ported to JXTA	.NET/My Services	public domain, supported by Sun Microsystems	de-facto common P2P platform
.NET / .NET My Services	proprietary & open-source standards	MIcrosoft Office and more	large base of MS customers	AOL, Sun J2EE/JXTA, Ximian MONO	MS internally	de-facto pervasive platform

Table 8. Business Comparison of Case Studies.

System/App Requirements	Type of the System		
	Centralized	**Client-Server**	**Peer-to-Peer**
decentralization	low (none)	high	very high
ad-hoc behavior	no	medium	high
cost of ownership	very high	high	low
anonymity	low (none)	medium	very high
scalability	low	high	high
performance	individual high / aggregate low	medium	individual low / aggregate high
fault resilience	individual high / aggregate low	medium	individual low / aggregate high
self-organization	medium	medium	medium
transparency	low	medium	medium
security	very high	high	low
interoperability	standardized	standardized	in progress

Table 9. Comparison of Solutions. *Darker shading represents more preferred characteristics or strengths.*

government funded (SETI@home), public domain efforts (Gnutella, JXTA), or privately owned companies (.NET My Services).

Finally, they have different business models, ranging from becoming an all-encompassing pervasive platform of choice (.NET My Services), or becoming a collaborative platform of choice (Groove, Magi), to selling more computers and using ads on screen savers (SETI@home) and becoming a P2P algorithm of choice.

7 LESSONS LEARNED

Peer-to-peer is a relatively new technology that is not yet proven in the research community and in industry. Therefore, it may be too early to make a firm statement on some of the lessons learned. Nevertheless, we can make some preliminary observations. Some of them are intuitive, while others are not obvious. We divide the observations into P2P strengths and weaknesses, P2P nontechnical challenges, and implications on users, developers, and IT.

7.1 Strengths and Weaknesses

A comparison of P2P with its alternatives, centralized and client-server systems, is summarized in Table 9. A more detailed comparison of P2P with its alternatives in specific classes of P2P systems is also provided in Appendix A. P2P systems are designed with the goals of decentralization, ad-hoc behavior, improved cost of

ownership, and anonymity. Therefore, P2P is superior to centralized and client-server systems in these areas. P2P has more decentralized control and data compared to alternatives; it supports systems whose parts can come and go and can communicate in an ad-hoc manner; the cost of ownership is distributed among the peers; and the peers can be anonymous. Compared to P2P, centralized systems are inherently centralized and client-server systems have centralized points of control and data at the servers.

It is harder to compare P2P with alternatives in terms of scalability, performance, security, self-organization, and fault-tolerance. It is our speculation that because of higher decentralization P2P can better satisfy these requirements as well. P2P has the potential to be more scalable than centralized and client-server solutions. However, P2P systems are highly dependent on the underlying topology and the types of applications. For example, mainframe computers are still competitive for transaction-based computing because they are optimized for this type of applications. On the other hand, P2P system are very suitable for large number of computers on the Internet, for mobile wireless users, or for device sensor networks. A similar reasoning applies to performance and fault-resilience. Individually, centralized systems have highly optimized performance and fault-tolerance, followed by client-server systems, and only then by P2P systems. The opposite is true for aggregate systems where P2P has higher aggregate performance and fault resilience by avoiding the need to interact with servers and having fewer points of failure.

It is unclear whether any of the systems has a clear advantage in self-organization and transparency. There has been long-standing work in improving these characteristics in both centralized systems and client-server systems. Self-organization is especially critical in Internet Data Centers, and mainframe computers also implement some forms of adaptation. Transparency has been the focus of client-server systems, in particular location transparency. While also a goal of P2P, transparency has not been sufficiently developed in P2P systems. Many actions require user intervention, such as connecting to a specific network of P2P systems, providing knowledge about the files and users, and understanding the failure model.

Finally, we believe that in two regards P2P lags behind its alternatives. Inherently, P2P systems expose more security threats than centralized and client-server models. Because of the nature of peer-to-peer interaction, sufficient guarantees or additional trust has to be established among the peers. Interoperability is matter of investment

and as development proceeds more interoperability may evolve. At the time of writing, client-server systems support the most interoperability.

7.2 Non-Technical Challenges

In addition to all of the technical motivations and difficulties with developing P2P systems, there are also non-technical challenges to the success of P2P. In spite of excellent technology, if these challenges are not overcome, P2P will likely remain an approach used only in small niches or by specific communities.

The chief challenge of P2P systems is *acceptance and use*. Because peers rely on one another to provide service, it is essential that numerous peers be available for the service to be useful. By comparison, centralized or client-server environments are potentially useful as long as the service provider keeps the service running. This is not the case in P2P systems. If the peers abandon the system, there are no services available to anyone.

P2P systems live and die on their network effects, which draw in more and more users. The value of network effects was informally specified by Metcalfe's Law, which states that "the utility of a system grows with the square of the number of users." While we cannot truly quantify utility to prove the law, it resonates clearly with the idea that more users make a system more useful.

Studies have also shown that all users are not alike, and that some users can actually damage a system. [Adar and Huberman2000] showed that in the Gnutella system, many users download files, but few actual provide files. Further, those with poor resources, such as bandwidth, can actually slow the network down by becoming bottlenecks for the system as a whole. This situation has been likened to the tragedy of the commons where poorly behaving users of a free, shared resource make the resource unusable by all.

Solutions to this problem revolve around building an economy around the use of the shared resource. One example of this is MojoNation [MojoNation 2001]. In MojoNation, users accumulate a form of currency called "mojo" by providing service to others. In the MojoNation scheme, this currency would be redeemable for other services or potentially from real-world vendors in the form of gift certificates or related benefits. However, implementing a MojoNation type scheme requires accounting to be performed, and this can limit anonymity and other potential benefits of P2P systems.

Related to acceptance and use is the *danger of fragmentation* of the user base. Typically, individuals are only going to participate in one or a very few different P2P systems because they simply don't have the resources to support multiple systems at the same time. Because of this, as each new system is introduced it fragments the user base and can potentially damage the value of all P2P systems. Consider the success of Napster. It was arguably not the most technically sound P2P system developed, but it was the first to gather a large user base. Therefore, until it was forcibly shutdown, it thrived because the network effects continued to draw more and more users.

Since Napster's demise, a number of music sharing systems, such as Gnutella and KaZaa have been developed and are both free and technically sound. However, neither has yet become the one truly dominant system, so neither has become as truly useful as Napster was in its prime.

Instant messaging faces a similar difficulty. While there are numerous IM systems, each with a significant user base, the overall utility of IM is limited because of fragmentation. Typically, to communicate with everyone a user wishes to reach, the user must maintain accounts and run clients from many IM systems at the same time. Interoperability, as discussed in Section 4.11, seems to be the only solution to this problem.

While P2P systems rely on scale for success, *scale* is also a significant challenge. In centralized systems, problems related to scale are relatively well understood and solutions are pretty well known. In the worst case, bigger, faster computing platforms are an option for improving scale. Decentralized P2P systems more often require algorithmic solutions to problems of scale. There is no central place to simply throw more computing resources. These distributed algorithms tend to be some of the most difficult to develop because they require decisions to be made at each local peer, usually with little global knowledge.

A final challenge involves the *release of control*. In centralized systems, the operator of the system has some control of the system. For example, an operator can monitor each transaction and potentially determine who initiated the transaction and how long it took. In a decentralized system, no such monitoring is possible. This scares many providers of traditional services so they resist P2P schemes. Certainly, the release of control can be seen as part of the reason for the Recording Industry Association of America's (RIAA) lawsuit against Napster. Napster provides an alternative method for music distribution. While copyright, payment, and fair-use issues are certainly at the center of the case, embracing a

Target	Criteria	Type of the System		
		Centralized	Client-Server	P2P
User	**Pervasiveness**	low	medium	high
	State-of-the-art	low	high	medium
	Complexity	high	low	medium
	Trust & Reputation	high	medium	low
Developer	**Complexity**	high	straightforward	typical - no atypical - yes
	Sustainability	low	high	medium
	Tools	medium (proprietary)	high (standardized)	low (few tools)
	Compatibility	medium	high	low
IT	**Accountability**	high	medium	low
	Being in control	high (fully)	medium	low
	Manageability	medium	high	low
	Standards	medium (proprietary)	high	low (inexistent)

Table 10. Comparison of Solutions. *Darker shading represents preferred characteristics.*

P2P approach to distribution, such as Napster, also implies a release of control of the distribution channel. This may explain why the RIAA and Napster have not been able to reach an agreement.

7.3 Implications for Users, Developers, and IT

P2P is not a solution for every future system and application. As we have seen in the previous two sections, it has both strengths and weaknesses. In this section, we evaluate the implications that P2P has for users, developers, and IT department. In Table 10, we compare P2P with its alternatives. P2P has the following implications for the users of P2P systems and applications: pervasiveness, complexity of use, state of the art, and trust and reputation.

Pervasiveness is becoming more and more important. While it may not be possible to access traditional services at any point on Earth at any time, P2P offers opportunities in this regard by relying on peers to provide services when there is no other infrastructure or means of access available. For example, cellular connectivity may not be available to the server, but a peer can offer the service locally.

Traditional centralized solutions have complex support compared to client-server solutions. P2P solutions are not as simple or as well-understood as the client-server model, yet wide deployment and use of Napster, Gnutel-la-based solutions, and other startups offers potential in this regard.

Currently, the client-server model represents the state of the art, but there is a lot of ongoing development in P2P research – second generation P2P systems (CAN, Pastry, CHORDS, Tapestry, etc.); open-source – JXTA; proprietary systems – .NET; and standards – P2PWG.

From the user perspective, P2P is probably weakest in the sense trust and reputation. Owners want to have trust in the service they are using, so they will use only reputable sites unless price is the main objective or the service is free. Centralized systems and the client-server model have traditionally built up trust and reputation, and this is a concern for users of P2P.

Developers are also concerned with complexity, as well as with the sustainability of solutions, with the availability of the tools, and the compatibility of the systems. In this comparison, the client-server dominate over the other solutions.

Client-server systems are well-understood and documented for developers. P2P systems offer promise in simple cases, such as document exchange, but for more complex requirements, such as collaborative applications, they require more complex algorithms and understanding. It is a similar case of sustainability. In the long term, centralized solutions may not be as sustainable as peer-to-peer solutions, but the client-server model will continue to be supported. Finally, the weakest aspect of P2P is the lack of tools and compatibility across various P2P systems.

P2P has the following implications for IT: accountability, being in control, manageability, and standards. The first three are very closely tied. Accountability is emphasized in centralized systems where access is monitored through logins, accounts, and the logging of activities. Accountability is more difficult to achieve in client-server systems, because of interactions with multiple clients. It is weakest in P2P systems, because of equal rights and functionality among the peers. Similar reasoning applies for being in control. In centralized and client-server systems, control is exercised at one or more well-defined points, whereas it is harder to achieve in P2P systems, where control is entirely distributed.

A similar situation exists for manageability. There are a lot of tools for the management of client-server systems, somewhat fewer for centralized systems, and fewest for P2P. Similar applies for standards. Most standards have been developed for client-server systems and very few for P2P systems.

While Table 9. compares the potential of P2P versus its alternatives in terms of the characteristics, Table 10. summarizes the existing implications of P2P on users, developers, and IT. In summary, there is a lot of potential for P2P, but it has not yet been realized.

8 SUMMARY AND FUTURE WORK

In this paper, we surveyed the field of P2P systems. We defined the field through terminology, architectures, goals, components, and challenges. We also introduced taxonomies for P2P systems, applications, and markets. Based on this information, we summarized P2P system characteristics. Then, we surveyed different P2P system categories, as well as P2P markets. Out of systems presented in Section 5, we selected eight case studies and described them in more detail. We also compared them based on the characteristics we introduced in Section 4. Based on this information, we derived some lessons about P2P applications and systems.

In the rest of this section, we revisit what P2P is, we explain why we think that P2P is an important technology, and finally we present the outlook for the P2P future.

8.1 Final Thoughts on What P2P Is

One of the most contentious aspects in writing this paper was to define what P2P is and what it is not. Even after completing this effort, we do not feel compelled to offer a concise definition and a recipe of what P2P is and what it is not. A simple answer is that P2P is many things to many people and it is not possible to come up with a simplified answer. P2P is a mind set, a model, an implementation choice, and property of a system or an environment.

- **A mind set.** As a mind set, P2P is a system and/or application that either (1) takes advantage of resources at the edge of the system or (2) supports direct interaction among its users. Such a system and/or application remains P2P regardless of its model or implementation. Examples include SETI@home, which is considered to have a client-server model, but displays the first mind set property, and Slashdot [Slashdot 2002], which enables the second mind set property, but really has a centralized implementation.
- **A model.** A system and/or application supporting the model presented in Figure 1 is P2P. In its purest form, P2P is represented by Gnutella, Freenet, and Groove. According to this, SETI@home does not have a P2P model, whereas Napster has a hybrid model.
- **An implementation choice.** P2P systems and applications can be implemented in a P2P way, such as JXTA

or Magi. However, a non-P2P application can also be implemented in a P2P way. For example, application-layer multicast can have a P2P implementation, and parts of the ORBs or DNS servers are also implemented in a P2P way.
- **A property of a system or an environment.** .NET as well as environments, such as small device sensor networks, may require P2P implementation solutions while not necessarily supporting a P2P application. P2P solutions may be required for scalability, performance, or simply because of the lack of any kind of infrastructure, making P2P the only way to communicate. This is similar to looking at P2P as an implementation choice, however in this case P2P is the *forced* implementation choice.

8.2 Why We Think P2P is Important

As P2P becomes more mature, its future infrastructures will improve. There will be increased interoperability, more connections to the (Internet) world, and more robust software and hardware. Nevertheless, some inherent problems will remain. P2P will remain an important approach for the following reasons.

- Scalability will always be a problem at certain levels (network, system, and application), especially with global connectivity, much of it wireless. It will be hard to predict and guarantee all service-level agreements. P2P can contribute to each area.
- Certain parts of the world will not be covered by (sufficient) connectivity, requiring ad-hoc, decentralized groups to be formed. P2P is a well-suited alternative when there is a lack of infrastructure.
- Certain configurations of systems and applications will inherently be P2P and will lend themselves to P2P solutions.

8.3 P2P in the Future

The authors of this paper believe that there are at least three ways in which P2P may have impact in the future:

- **P2P algorithms** probably have the biggest chance of making impact. As the world becomes increasingly decentralized and connected, there will be a growing need for P2P algorithms to overcome the scalability, anonymity, and connectivity problems.
- **P2P applications** are the next most likely to succeed in the future. Examples, such as Napster are a convincing proof of such a possibility.
- **P2P platforms** are the third possible scenario for P2P. Platforms such as JXTA may be widely adopted, in

which case many other P2P systems can also gain wide adoption.

8.4 Summary

We believe that P2P is an important technology that has already found its way into existing products and research projects. It will remain an important solution to certain inherent problems in distributed systems. It may not be the only solution and may not be appropriate for all problems, but it will continue to be a strong alternative for scalability, anonymity, and fault resilience requirements. P2P algorithms, applications, and platforms have an opportunity for deployment in the future. From the market perspective, cost of ownership may be the driving factor for P2P. The strong presence of P2P products indicates that P2P is not only an interesting research technology but also a promising product base.

REFERENCES

ABERER, K., PUNCEVA, M., HAUSWIRTH, M., SCHMIDT, R. 2002. Improving Data Access in P2P Systems. IEEE Internet Computing 6(1):58-67. January-February.

ADAMIC, L. The Small World Web. 2000. *Technical Report,* Xerox Palo Alto Research Center.

ADAR, E. AND HUBERMAN, B. 2000. Free Riding on Gnutella. *First Monday, vol 5, no 10 (October).* (see also www.firstmonday.dk/issues/issue5_10/adar.)

AKAMAI 2001. www.akamai.com.

ALBITZ, P, AND LIU. C. 1995. DNS and BIND. O'Reilly and Associates.

ANDY O. 2001. Peer-To-Peer, Harnessing the Power of Disruptive Technology. O'Reilly.

AOL 2001. www.aol.com.

APPLIED METACOMPUTING. 2000. Legion - An Integrated Architecture for Secure Resource Sharing. Applied Meta Computing *Peer-to-Peer Architectural Proposal.*

ARIBA, IBM, AND MICROSOFT. 2000. UDDI Technical White Paper, available at www.uddi.org, September 2000.

AVAKI CORPORATION. 2001. Avaki 2.0 Concepts and Architecture. White paper. www.avaki.com/papers/AVAKI_concepts_architecture.pdf.

BEARSHARE 2001. www.bearshare.com.

BALL, T. AND RAJAMANI, S. K. (2001) Automatically validating temporal safety properties of interfaces. In the Proceedings of the International SPIN Workshop on Model Checking of Software, pp. 333-344, May 2001.

BARAK, A. AND LITMAN, A. 1985. MOS: a Multicomputer Distributed Operating System. *Software - Practice and Experience*, 15(8):725–737. August.

BARAK, A. AND WHEELER, R. 1989. MOSIX: An Integrated Multiprocessor UNIX. *Proceedings of the Winter 1989 USENIX Conference*, pages 101–112. February.

BECKER D.J., STERLING T., SAVARESE D., DORBAND J.E., RANAWAK U.A., PACKER C.V. 1995. "Beowulf: A Parallel Workstation for Scientific Computation", Proceedings of the International Conference on Parallel Processing.

BERNERS-LEE T. FIELDING, R., MASINTER, L. 1998. Uniform Resource Identifiers (URI): Generic Syntax. IETF Request for Comments. August 1998

BERNSTEIN, P. A. 1996. Middleware: A Model for Distributed System Services. *Communications of the ACM*, 39(2):86–98.

BOLCER, G. ET AL. 2000. Peer-to-Peer Architectures and the Magi Open Source Infrastructure. *Endeavors technologies White Paper.*

BOLCER, G, 2001. Magi: An Architecture for Mobile and Disconnected Workflow. *White paper* (www.endeavors.com)

BOLOSKY, W., DOUCEUR, J., ELY, D., AND THEIMER, M. 2000. Feasibility of a Serverless Distributed File System Deployed on an Existing Set of Desktop PCs. *Proceedings of SIGMETRICS,* Santa Clara, CA, USA, Jun, 2000.

BOND, J. 2001. Business Uses of Peer to Peer (P2P) Technologies. Netmarkets Europe *White Paper.* (www.netmarketseurope.com).

BOX, D. EHNEBUSKE, D., KAKIVAYA, G., LAYMAN, A. MENDELSOHN, N., NIELSEN, H.F., THATTE, S., WINER. D. 2000. Simple Object Access Protocol (SOAP) 1.1. W3C Note 08 May 2000.

BRAY, T., PAOLI, J., SPERBERG-MCQUEEN, C. M., MALER, E. (Editors). 2000. Extensible Markup Language (XML) 1.0 (Second Edition) W3C Recommendation 6 October 2000.

BRITTON, C. 2000. IT Architectures and Middleware: Strategies for Building Large, Integrated Systems. Addison Wesley.

BUYYA, R. 2001. Economic models for Management of the Resources in Peer-to-Peer and Grid Computing. *Proceedings of the Commercial Applications for High-Performance Computing Conference.*

CARPENTER. B. 2000. Internet Transparency. Internet Society Request for Comments 2775. www.faqs.org/rfcs/rfc2775.html.

CASAVANT, T. L. AND KUHL, J. 1988. A Taxonomy of Scheduling in General-Purpose Distributed Computing Systems. *IEEE Transactions on Software Engineering*, SE-14(2):141–152.

CASTRO, M. AND LISKOV, B. 1999. Practical Byzantine Fault Tolerance. *Proc. Usenix Symp. Operating Systems Design and Implementation* (OSDI 99) Usenix Assoc., Berkeley, California.

CHAPIN, S. J. 1996. Distributed and Multiprocessor Scheduling. *ACM Computing Surveys*, 28(1):233–235.

CHAUM, D. 1981. Untraceable Electronic Mail Return Addresses and Digital Pseudonyms, Communication of the ACM 24, 2, Feb. 1981, pp. 84-88.

CHILDERS, L., DISZ, T., OLSON, R, PAPKA, M., STEVENS, R., AND UDESHI, T. 2000. Access Grid: Immersive Group-to-Group Collaborative Visualization. *Pro-

ceedings of the Fourth International Immersive Projection Technology Workshop. Jun, 2000. pp?

CHRISTENSEN, C.M. 1997. The Innovator's Dilemma. Harvard Business School Press.

CHRISTENSEN, E., CURBERA, F., MEREDITH, G., WEERAWARANA, S. 2001. Web Services Description Language (WSDL) 1.1, W3C Note 15 March 2001. www.w3.org/TR/wsdl.

CLARKE, I. 1999. A Distributed Decentralized Information Storage and Retrieval System, unpublished report, Division of Informatics, University of Edinburgh. freenet.sourceforge.net/freenet.pdf.

CLARKE, I., I., MILLER, S.G. WONG, T.W., SANDBERG, O., WILEY, B. 2002. PROTECTING FREE EXPRESSION ONLINE WITH Freenet. IEEE Internet Computing 6(1):40-49. January-February.

CLARKE, I., SANDBERG, O, WILEY, B., HONG, T.W. 2001. Freenet: A Distributed Anonymous Information Storage and Retrieval System. *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability,* LNCS 2009, ed. by H. Federrath. Springer: New York.

CLIP 2 MARKETING 2000. Gnutella: To the Bandwidth Barrier and Beyond. www.clip2.com/gnutella.html.

COULORIS, G., DOLLIMORE, J. 2001. Distributed Systems. Concepts and Design. Addison Wesley.

DABEK, F., BRUNSKILL, E. KAASHOEK, F., KARGER, D., MORRIS, R., STOICA, I., AND BALAKRISHNAN, H. 2001. Building Peer-to-Peer Systems With Chord, a Distributed Lookup Service. *Proceedings of the 8th Workshop on Hot Topics in Operating Systems (HotOS-VIII),* Schloss Elmau, Germany, May 2001.

DATASYNAPSE 2001. The Next Generation of High Performance Computing For Financial Services. White Paper.

DENNING. D. E. 1976. A Lattice Model of Secure Information Flow. Communications of the ACM 19(5):236-243, May 1976.

DRUSCHEL P. AND ROWSTRON, A. 2001. PAST: A Large-Scale, Persistent Peer-to-Peer Storage Utility, HotOS VIII, Schloss Elmau, Germany, May 2001.

EDWARDS, W., MYNATT, T., PETERSEN, K., SPREITZER, M., TERRY, D., THEIMER, M. 1997. Designing and Implementing Asynchronous Collaborative Applications with Bayou. Proceedings of the Tenth ACM Symposium on User Interface Software and Technology (UIST), Banff, Alberta, Canada, October, 1997.

ENDEAVORS TECHNOLOGY 2001. www.endeavors.com.

EXODUS 2001. www.exodus.com

FASTTRACK. 2001. Product Description. www.fast-track.nu/index_int.html.

FATTAH, H.M., FATTAH, H.M. 2002. P2P: How Peer-to-Peer Technology Is Revolutionizing the Way We Do Business, Dearborn Trade

FOSTER, I. 2000. Internet Computing and the emerging Grid. Nature, Dec. 7.

FOSTER, I., KESSELMAN, C. 1999. The Grid: Blueprint for a New Computing Infrastructure. Morgan Kaufman Publishers, Inc. San Francisco, California.

FOLDING@HOME 2001. foldingathome.stanford.edu.

FREENET. 2001. The FreeNet home page, freenet.sourceforge.net. www.freenetproject.org.

GABBER, E., GIBBONS, P., KRISTOL, D., MATIAS, Y., AND MAYER, A. 1999. Consistent, yet anonymous, Web access with LPWA. Communications of the ACM. 42, 2. February 1999. pp. 42-47.

GARTNER CONSULTING 2001. The Emergence of Distributed Content Management and Peer-to-Peer Content Networks Engagement. *Report* #010022501.

GENOME@HOME. 2001. genomeathome.stanford.edu.

GNUTELLA. 2001. The Gnutella home page, gnutella.wego.com.

GOSCINSKI, A. 1991. Distributed Operating Systems. The Logical Design. Addison Wesley Publishing.

GRAHAM, R.L. 2001. Traditional and Non-Traditional Applications; Peer-to-Peer Networks. *Lecture.* www.ida.liu.se/~TDTS43/tdts43-10-peer-to-peer.pdf.

GRIBBLE, S., HALEVY, A., IVES, Z., RODRIG, M., AND SUCIU, D. 2001. What Can Peer-to-Peer Do for Databases and Vice Versa? *Proceedings of the WebDB: Workshop on Databases and the Web, Santa Barbara, CA, USA.*

GRID COMPUTING BIOTECHNOLOGY. 2001. www.gridcomputing.net.

GRID FORUM. 2001. www.gridforum.org.

GRIMSHAW, A., WULF, W.A., AND THE LEGION TEAM. 1997. The Legion Vision of a Worldwide Virtual Computer. *Communications of the ACM, 40(1):39-45.*

GRIMSHAW, A., Easy to use Object-oriented Parallel Programming with mentat. *IEEE Computer, pp39-51, May 1993.*

GRIMSHAW, A. S., WULF, W. A., FRENCH, J. C., WEAVER, A. C., REYNOLDS JR., P. F. 1994. Legion: The Next Logical Step Toward a Nation-wide Virtual Computer. UVa CS Technical Report CS-94-21, June 8, 1994.

44

GROOVE NETWORKS. 2000. Introduction to Groove. *Groove Networks White Paper*.

GROOVE NETWORKS. 2000. Connection Age. *Groove Networks White Paper*.

GROOVE NETWORKS. 2001. Groove Networks Product Backgrounder. *Groove Networks White Paper*. www.groove.net/pdf/ groove_product_backgrounder.pdf.

GROOVE NETWORKS. 2001. Groove and .NET Brief. *Groove Networks Brief*. www.groove.net/feature/ webandpeer/groove_and_dotnet.pdf.

HAC, A. 1989. Load Balancing in Distributed Systems: A Summary. *Performance Evaluation Review*, 16:17–25.

HALL GAILEY, J. 2001. Introducing .NET My Services. MSDN Library note, msdn.microsoft.com/library/ default.asp?url=/library/en-us/dndotnet/html/myser-vintro.asp. September 2001.

HEYLIGHEN, F. 1997. Principa Cybernetica Web. pespmc1.vub.ac.be/SELFORG.html.

HOUSTON, B. 2001. The "GnutellaMandragore" virus. www.exocortex.org/gnutella.

HOWARD, J.H., KAZAR, M.L., MENEES, S.G., NICHOLS, D.A., SATYANARAYANAN, M., SIDEBOTHAM, R.N., WEST, M.J. 1988. Scale and Performance in a Distributed File System. *ACM Transactions on Computer Systems, Feb. 1988, Vol. 6, No. 1, pp. 51-81*.

HOWES, T., SMITH, M. 1997. LDAP. Programming Directory-Enabled Applications with Lightweight Directory Access Protocol. MacMillan Technical Publishing, USA.

ICQ 2001. web.icq.com.

IEEE 1990. Institute of Electrical and Electronics Engineers. IEEE Standard Computer Dictionary: A Compilation of IEEE Standard Computer Glossaries. New York, NY: 1990.

IEEE 2001. www.ieee.org.

INTEL 2001. Peer-to-Peer-Enabled Distributed Computing. *Intel White Paper*.

CARPENTER. B. 2000. Internet Transparency. Internet Society Request for Comments 2775. www.faqs.org/ rfcs/rfc2775.html.

JXTA 2001.The JXTA home page www.jxta.org.

KASREL, B. 2001. P2P's Pervasive Future. *The Forrester Report*.

KATZENBEISSER, S. 1999. Information hiding techniques for steganography and digital watermarking. *Artech House Books*.

KaZaA. www.kazaa.com. 2001.

KINDBERG, T. 2002. Personal Communication.

KOTZEN, M. 2001. Dramatic Improvements in the Gnutella Network Since Late 2000. www.limewire.com/ index.jsp/net_improvements.

KUBIATOWICZ, J., BINDEL, D., CHEN, Y., CZERWINSKI, S., EATON, P., GEELS, D., GUMMADI, R., RHEA, R., WEATHERSPOON, H., WEIMER, W., WELLS, C., AND ZHAO, B. 2000. OceanStore: An Architecture for Global-Scale Persistent Storage. *Proceedings of the Ninth international Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2000),* November 2000.

LEIGH, P. AND BENYOLA, P., 2001. Future Developments in Peer Networking. Equity Research, *White Paper,* Raymond James & Associates, INC.

LIMEWIRE 2001. www.limewire.com.

LITZKOW, M. AND SOLOMON, M. 1992. Supporting Checkpointing and Process Migration outside the UNIX Kernel. *Proceedings of the USENIX Winter Conference, pages 283–290.*

LITZKOW, M., LIVNY, M., AND MUTKA, M. June 1988. Condor - A Hunter of Idle Workstations. *Proceedings of the 8th International Conference on Distributed Computing Systems*, pages 104–111.

MICROSOFT. 2001. Microsoft .NET Passport Technical Overview. September 2001.

MICROSOFT. 2001a. Microsoft Windows 2000 Data Sheet. www.microsoft.com/education/ ?ID=Win2kDatasheet.

MILGRAM, S. 1967. The Small World Problem. Psychology today. 1, 61.

MILLER, M. 2001. Discovering P2P. Sybex, to appear.

MOCKAPETRIS, P. 1989. DNS Encoding of Network Names and Other Types. *Internet Request For Comment 1101* Network Information Center, SRI International, Menlo Park, California.

MOJONATION. 2001. http://mojonation.net.

MORE, M. 2001. P2P: Getting Down to Business. *Many-Worlds White Paper*. www. manyworlds.com.

MOORE, D., HEBELER, J. 2001. Peer-to-Peer: Building Secure, Scalable, and Manageable Networks. McGraw Hill.

MORGAN, J. 2002. Personal Communication.

NAPSTER. 2001. The Napster home page, www.napster.com.

NATARAJAN. A., ET AL, Studying Protein Folding on the Grid: Experiences Using CHARMM on NPACI under Legion. *High Performance Distributed Computing 10, August 7-9, 2001*

NECULA, G. 1997. Proof-Carrying Code. Proceedings of the 24th Annual ACM Symposium on Principles of Programming Languages, pp 106-119. Paris France, January 1997.

NECULA, G., AND LEE, P. 1998. The Design and Implementation of a Certifying Compiler. Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation, pp 106-119. Montreal, Canada, pp. 333-344, June 1998.

NETMEETING BY MICROSOFT 2001. www.microsoft.com/windows/netmeeting/default.asp.

NETHISINGHE, S., 2001. TV or the PC? Consumer Perspective. Philips Presentation.

NOWITZ, D. 1978. UUCP Implementation Description. UNIX Programmer's Manual, Bell Laboratories, October, 1978.

OBJECT MANAGEMENT GROUP (OMG) 2001. www.omg.org.

OMG. 1996. Common Object Request Broker Architecture and Specification. *Object Management Group Document Number 96.03.04.*

OPENCOLA 2001. www.opencola.com

OZZIE, J. 2000. SYSTEMS INTEGRATION WITH GROOVE. *Groove Networks Presentation.* devzone.groove.net/library/Presentations/GrooveSystemsIntegratoin.ppt.

PERKINS, C., 2001. Ad Hoc Networking. Addison Wesley.

PEER-TO-PEER WORKING GROUP. 2001. www.p2pwg.org.

PEER-TO-PEER WORKING GROUP. 2001. Bidirectional Peer-to-Peer Communication with Interposing Firewalls and NATs. p2pwg White Paper, Revision 0.091. May 23, 2001. http://www.peer-to-peer-wg.org/tech/nat/.

PFITZMANN, A., WAIDNER, M. 1987. Networks without User Observability. Computer Security 2, 6, pp.158-166.

PLATFORM COMPUTING. 2001. www.platform.com.

RAMANATHAN, M. K., KALOGERAKI, V. PRUYNE, J. 2001. Finding Good Peers in the Peer-to-Peer Networks. *HPL Technical Report HPL-2001-271.*

RATNER, D. 1998. Roam: A Scalable Replicated System for Mobile and Distributed Computing. *UCLA Technical Report UCLA-CSD-970044.*

RATNASAMY, S., FRANCIS, P., HANDLEY, M., KARP, R., SHENKER, S. 2001. A Scalable Content-Addressable Network. *Proceedings of the SIGCOMM, pp 161-172.*

REITER, M. K., RUBIN, A. D. 1998. Crowss: Anonymity for Web Transactions. ACM Transaction on Information and System Security 1, 1, November 1998, pp 66-92.

RHEA, S., WELLS, C., EATON, P., GEELS, D., ZHAO, B., WEATHERSPOON, H., AND KUBIATOWICZ, J. 2001. Maintenance-Free Global Data Storage. Internet Computing, vol. vol 5, no 4. pp 40-49. September 2001

ROSENBERRY, W., KENNEY, D., AND FISHER, G. 1992. Understanding DCE. O'Reilly & Associates, Inc.

ROMAN, G., HUANG, Q., HAZEMI, A. 2001. Consistent Group Membership in Ad Hoc Networks. Proceedings of ICSE, Toronto, Canada.

ROWSTRON, A. AND DRUSCHEL, P. 2001. Storage Management and Caching in PAST, a Large-Scale, Persistent, Peer-to-Peer Storage Utility. *Proceedings of SOSP,* pp 188-201.

SALTZER, J.H., REED, D.P., CLARK, D.D. 1984. End-To-End Arguments in System Design, ACM TOCS, Vol 2, Number 4, November 1984, pp 277-288.

SANDBERG, R., D. GOLDBERG, S. KLEIMAN, D. WALSH, B. LYON. 1985. Design and Implementation of the Sun Network Filesystem. *USENIX Conference Proceedings,* USENIX Association, Berkeley, CA, Summer 1985.

SAROIU, S., GUMMADI, P., AND GRIBBLE, S. 2002. A Measurement Study of Peer-to-Peer File Sharing Systems. MMCN, San Jose, CA, USA January 2002.

SATYANARAYANAN, M. 2001. Pervasive Computing: Vision and Challenges. *IEEE Personal Communications, August 2001.*

SATYANARAYANAN, M., KISTLER, J.J., KUMAR, P., OKASAKI, M.E., SIEGEL, E.H., AND STEERE, D.C. 1990. Coda: A Highly Available File System for a Distributed Workstation Environment. IEEE Transactions on Computers, pages 447-459,vol 39 1990.

SCARLATA, V., LEVINE, B. N., SHIELDS, C. 2001. Responder Anonymity and Anonymous Peer-to-Peer File Sharing. ICNP'01.

SCHWARTZ, E. 2001. 'Parasitic Grid' Wireless Movement May Threaten Telecom Profits. August 24, 2001. http://www.infoworld.com/articles/hn/xml/01/08/24/010824hnfreewireless.xml?0827mnam.

SHA-1. 1997. American National Standards Institute, American National Standard X9.30.2-1997: Public Key Cryptography for the Financial Services Industry - Part 2: The Secure Hash Algorithm (SHA-1).

SHAMIR, A. 1979. How to share a secret. Communications of the ACM, vol. 22, n. 11, pp. 612-613, Nov. 1979.

SHILEDS, C., LEVINE, B. N. 2000. A protocol for Anony-

mous Communication Over the Internet. 7th ACM Conference on Computer and Communication Security (ACM CCS 2000), November 2000.

SHIRKY, C. 2001. What is P2P... and what Isn't. *An article published on O'Reilly Network.* www.openp2p.com/lpt/a//p2p/2000/11/24/shirky1-whatisp2p.html.

SEARCHLING. 2000. SearchLing Peer-to-Peer Networking. White Paper. October 2000.

SETI@HOME 2001. setiathome.ssl.berkeley.edu.

SINGH, M., 2001. Peering at Peer-to-Peer Computing. *IEEE Internet Computing, 5(1):4-5.*

SINGH, M., YU, B., VENKATRAMAN, M. 2001. Community-Based Service Location. *Communications of the ACM, 44(4):49-54.*

SOFTWAX 2001. SoftWax Technologies. www.softwax.com.

STOICA, I., MORRIS, R., KARGER, D., KAASHOEK, F., AND BALAKRISHNAN, H. 2001. Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications. *Proceedings of the SIGCOMM, pp 149-160.*

STROM, D. 2001. Businesses Embrace Instant Messaging. January 2001 enterprise.cnet.com/enterprise/0-9534-7-4403317.html.

SULLIVAN III, W. T., WERTHIMER, D., BOWYER, S., COBB, J., GEDYE, D., ANDERSON, D. 1997. A new major SETI project based on Project Serendip data and 100,000 personal computers. Astronomical and Biochemical Origins and the Search for Life in the Universe. Proc. of the Fifth Intl. Conf. on Bioastronomy, Colloq. No. 161, eds. C.B. Cosmovici, S. Bowyer, and D. Werthimer (Publisher: Editrice Compositori, Bologna, Italy.

SUTHAR, P. AND OZZIE, J. 2000. The Groove Platform Architecture. *Groove Networks Presentation.* devzone.groove.net/library/Presentations/GrooveApplicationArchitecture.ppt.

SYVERSON, P. F., GOLDSCHLAG, D. M., REED M. G. 1997. Anonymous Connections and Onion Routing, 1997 IEEE Symposium on Security and Privacy. pp. 44-53.

TANENBAUM, A. S. 1981. Computer Networks, Prentice-Hall International.

TANENBAUM, A. S. 1992. Modern Operating Systems. Prentice Hall International.

TEXAR CORPORATION. 2001. Peer-to-Peer Computing: Issues and Opportunities in Information Sharing. *A White Paper.*

TOADNODE 2001. www.toadnode.com.

VARELA, F., H. MATURANA, AND R. URIBE, 1974. Auto-

poiesis: The Organization of Living Systems, Its Characterization and a Model. *_Biosystems_, Vol. 5 (1974), pp. 187-196.* Amsterdam: North-Holland.

VEYTSEL, A. 2001. There is no P-to-P Market... But There is a Market for P-to-P. *Aberdeen Group Presentation* at the P2PWG, May 2001.

XDEGREES 2001. www.xdegrees.com.

W3C Note 08 May 2000

WALDMAN, M., RUBIN, A. AND CRANOR, L 2000. Publius: A Robust, Tamper-Evident, Censorship-Resistant Web Publishing System. Proceedings of the USENIX Security Symposium. Denver, Colorado, USA. Aug, 2000.

WALDO, J., WYANT, G., WOLLRATH, A., KENDALL, S. 1997. "A Note on Distributed Computing", Mobile Object Systems, Lecture Notes in Computer Science, No. 1222, pp. 49-64, Springer-Verlag, Berlin (D).

WANG, Y.-T. AND MORRIS, R. J. T. 1985. Load Sharing in Distributed Systems. *IEEE Transactions on Computers*, C-34(3):204–217.

WATERHOUSE, S., DOOLIN, D.M., KAN G., FAYBISHENKO, Y. 2002. Distributed Search in P2P Networks. IEEE Internet Computing 6(1):68-72. January-February.

WATTS, D.J. The Dynamics of Networks between Order and Randomness. Princeton University Press.

WATTS D.J. AND S. STROGATZ, H. 1998. Collective dynamics of `small-world' networks, *Nature 393, 440--442.*

WILLIAMS RICE, A. AND MAHON B. 2000. Peer Networking. *Deutsche Banc Alex. Brown White Paper.*

WOLLRATH, A., ET AL. 1996. A Distributed Object Model for the Java System. *Proceedings of the USENIX 1996 Conf. on Object-Oriented Technologies (COOTS)*, pp. 219-231.

XU, Z., MILLER, B., AND REPS, T., 2000. Safety Checking of Machine Code. *Proceedings of the SIGPLAN Conference on Programming Language Design and Implementation,* pp 70-82. Vancouver B.C., Canada. June 2000.

XU, Z., REPS, T., AND MILLER, B., 2001. Typestate Checking of Machine Code. *Proceedings of the 10th European Symposium on Programming,* pp 335-351. Genova, Italy, April 2001. Lecture Notes in Computer Science 2028, G. Goos, J. Hartmanis and J. van Leeuwen (Eds.)

YAGA 2001. Yaga, The Digital Market Place. www.yaga.com.

YAHOO! 2001. www.yahoo.com.

YANG, B; GARCIA-MOLINA, H. (2001) Comparing Hybrid Peer-to-Peer Systems. The VLDB Journal, pp 561-570, September 2001.

ZHAO, B., KUBIATOWICZ, J., JOSEPH, A. 2001. Tapestry: An Infrastructure for Fault-Tolerant Wide-area Location and Routing. Computer Science Division, University of California, Berkeley Technical Report No. UCB/CSD-01-1141, April 2001.

ZHOU, S., ZHENG, X., WANG, J., AND DELISLE, P. 1994. Utopia: A Load Sharing Facility for Large, Heterogeneous Distributed Computer Systems. *Software-Practice and Experience*, 23(2):1305-1336, December.

P2P Solutions	Comparison Criteria			
	Example Systems	Alternatives	Key Feature	Typical Target Architecture
Distributed Computing	Avaki, Entropia, DataSynapse, MojoNation, United Devices, cancer@home, SETI@home	grids, clusters, supercomputers	scalability, self-organization	personal/company owned computers
Data Sharing	Napster, Gnutella, Aimster, Freenet,	DFS, Web-based	availability	Internet
Collaboration	Magi, Groove Cybiko,	desktop, Web-based	ad-hoc nature	Internet, ad-hoc networks
Platforms	JXTA, .NET,	OS, Web, traditional distributed systems,	interoperability	Internet, enterprises

Table 11. The P2P Types of Systems.

APPENDIX A P2P VS. ALTERNATIVES

P2P is not a solution to every problem in the future of computing. Alternatives to P2P are traditional technologies, such as centralized systems and the client-server model. One potentially interesting observation is that these two alternatives to P2P are a hardware architecture and a programming model, which seems like a comparison of apples and oranges. This can be particularly awkward as we compare these solutions. However, P2P is indeed both an architecture (how we design and implement systems) and a programming model (how we program algorithms and in general conceive the systems and applications).

Centralized systems (possibly organized in small clusters) are represented by *mainframes* and *supercomputers* for execution of compute-intensive applications and large batch jobs, and *high-end servers* for data/content sharing and for collaborative applications. *Internet Data Centers* are also a form of centralized-point-of-control (even if distributed) installation of a large number of machines serving one purpose – content sharing.

The client-server model has different implementations, including on top of *clusters*, such as Platform Computing LSF, *wide-area middleware systems*, such as CORBA, RMI, and *Web-based solutions*, such as a Web server offering services to clients.

Some of the strengths of centralized and client-server systems include well-understood programming models, guaranteed performance, strong security, and well-known solutions for reliability. A lot of work has already been standardized and used for several decades. These systems are managed from a central point of control enabling good insight into the status and straightforward changes to configuration and parameters. On the negative side, these systems have limits to scalability, and a lot of them are proprietary and consequently hard to change. Centralized systems are costly to own and maintain and hard to deploy on a wide scale, such as in pervasive computing.

In Table 11, we summarize various types of P2P systems that will be compared in the rest of the section. We have classified P2P systems into those supporting distributed computing, data sharing, and collaboration, and into platforms. As alternatives, most of the P2P systems have Web-based tools, and some have traditional distributed systems. The key features of these systems are *scalability* by avoiding the bottlenecks of centralized solutions systems are better able to scale, *self-organization*, caused by individual peers continuously going up and down, the *ad-hoc nature* of peers that establish communication channels for collaboration, and interoperability as required by platforms.

In the following three tables, we compare each class of P2P system with its alternatives in more detail. In

Solutions	Comparison Criteria												
	Example	Infra-structure	Decentralization	Scalability (maximum)	Anonymity	Self-Organiz.	Cost of Ownership	Ad-Hoc	Performance	Security Threats	Transparency	Fault Resilience	Interoperability
Supercomputers High-End Servers	IBM Sysplex, Compaq, Sun, HP	standalone small LAN	somewhat	a few	low	low	high	low	high for small	minimal	low	elaborate at all levels	low
Clusters of PCs	Condor, LSF	Intranet, LANs	strong	a few 100	low	low	moderate	medium		low	high	checkpoint-restart fail-over	medium
P2P	Avaki, SETI@home	Internet intranets	significant	SETI (500000)	high	high	low (distributed)	high		high		restart on failure, tolerate disconnection	low

Table 12. Comparison of Distributed Computing Solutions.

49

Solutions	Comparison Criteria for Collaborative Computing Solutions													
	Example	Infra-structure	Conne-ctivity	Decentr alization	Scalability (maximum)	Anonym ity	Self-Organiz	Cost of Ownership	Ad-hoc	Perform ance	Security Threats	Transpa rency	Fault Resilience	Interope rability
Desktop-Based	Netmeeting, Lotus Notes	Intranet	limited	low	low	low	low	high	low	medium	low	low	low	medium
Web-Based	SharePoint, instant messaging (AOL, Yahoo)	Internet	Internet-wide	low	medium	medium	low	medium	low	high	high	low	high	high
P2P	Groove, Magi	Internet, ad-hoc, enterprise	World-wide	high	medium	medium	high	low	medium	low	very high	high	medium	low

Table 13. Comparison of Collaborative Computing Solutions.

Table 12, we compare P2P solutions for distributed computing with clusters of PCs, supercomputers and high-end servers, and grids. The earlier systems, such as high-end servers and supercomputers, run either standalone or in intranets. Examples are IBM mainframes (clusters of Sysplex machines), standalone or clusters of high-end UNIX servers. They have small scalability, elaborate fault-handling, and extensive security, but high cost of ownership. They are very expensive to obtain and to maintain.

Clusters represent a transition to P2P solutions in that they scale more compared to high-end systems, up to a few hundreds (e.g., LSF claims to have scaled to over 500 hundred machines), but they compromise in terms of security, fault handling, and cost of ownership. Similarly to high-end systems, they are deployed within intranets or even LANs.

Finally, P2P solutions are deployed on the Internet where they have increased scalability (e.g., up to 500,000 in the case of SETI@home), but they are very vulnerable to security attacks and have implemented only the coarsest failure tolerance such as restarting failed tasks. However, their cost of ownership is very small. It can be as little as the cost of an average PC.

Figure 25 shows how scalability increases in the transition from high-end solutions over clusters of computers to P2P solutions. On the same plot, the cost of ownership decreases. Simultaneously, the graphs of security threats are growing and fault tolerance are decreasing, but these
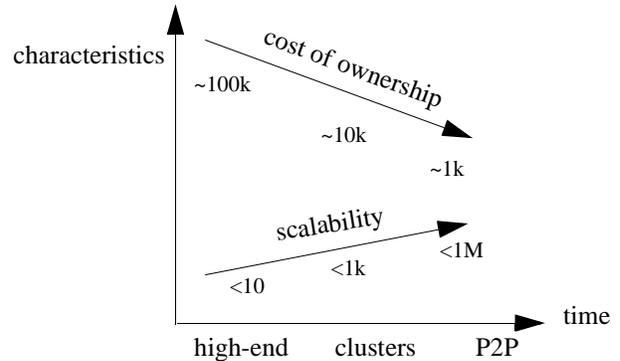

Figure 25: Comparison of Distributed Computing Solutions.
plots are omitted from the figure because of the lack of quantitative comparison numbers.

In Table 13, we compare collaborative P2P solutions with two other collaborative solutions: desktop-based computing in an intranet environment and Web-based collaborative solutions. Historically, there seems to be a progressive evolution of connectivity from intranets through the Internet to ad-hoc wireless connectivity. However, the security threats are increasing proportionally with the increase in connectivity.

In Table 14, we compare the solutions for context sharing with their historical alternatives: distributed file systems, and Web-based publishing. Similar to other comparisons, the cost of ownership is reduced, but security attacks are increased for P2P. In addition, consistency is decreased as well as availability, but anonymity is

Solutions	Comparison Criteria for Context Sharing														
	Example	Purpose	Infra-structure	Consistency	Decentra lization	Scalability (maximum)	Anonymity	Self-Organiz.	Cost of Owner ship	Ad-hoc	Perfor mance	Security Threats	Transpa rency	Fault Resilience	Interope rability
Distributed File Systems	NFS, AFS, DFS	general purpose file sharing	Clusters, WANs, Intranet	strong	medium	high	restricted	low	high	low	N/A	minimal	high	high	low
Web-Based	Web pages	information sharing	Internet	mostly read-only	medium	high	by obscu-rity	medium	low	low	N/A	high	medium	Internet-limited	high
P2P	Napster, Gnutella	content sharing	Internet, ad-hoc	weak	high	high	guaranteed	high	very low	high	N/A	very high	**low**	connectiv-ity-limited	**low**

Table 14. Comparison of Solutions for Context Sharing.

improved. Traditional file systems provided guarantees for acceptable information consistency at individual nodes in a Distributed File Systems (DFS), modulo NFS problems. Web based systems are largely considered read-only, even though there is increasingly more read-write content. This significantly simplifies consistency requirements. P2P systems cannot make significant guarantees for consistency.

Traditional DFSs are designed to be highly available and users are guaranteed access to the data by techniques, such as replication, caching, and failover. On the Web, caching also takes place, but network connectivity problems are the major concern for consistency, with obsolete caches as the next level of problems. In P2P systems, there is an assumption that there will be a best effort for availability. The content may or may not be available subject to connectivity and peer availability.

Finally, anonymity is restricted in a DFS. It is possible on the Web, largely by obscuring the identity. P2P systems bring it to the next level, by guaranteeing anonymity. For example, in systems like Freenet, Publius, or Free Haven, anonymity is guaranteed to the reader, writer, and publisher of the content.

In summary, we can compare P2P solutions with traditional client-server solutions based on the model (see Figure 4). Similarly, as in previous comparisons, the P2P solutions have improved scalability, fault resilience, anonymity, security, and self-organization (ad-hoc nature).