Chapter 4
Peer-to-Peer (P2P) computing

4.1 Introduction

In this chapter, P2P computing in general will come under discussion, as well as the definitions and concepts pertaining to this concept. The basic concept of P2P computing, however, is not new. The modern-day hype around and extensive media coverage of P2P computing can be attributed to the triggering and enabling of applications in the technological revolution of the 21st century.

The principal aim of this chapter is to introduce the reader to the concept of and the applications consequent upon P2P computing.

In so doing, the following sub-problems will come under discussion:

- What are the historical foundations at the core of the P2P computing revolution and evolution?
- Formulate a definition for and identify the aspects pertaining to P2P computing within the modern technological revolution.
- What are the existing and "music-of-the-future" applications of P2P computing?
- Does P2P computing technology, together with its applications, have a place in businesses and/or tertiary-education institutions, and can it be implemented in such environments?

The first half of this chapter will be used to take a closer look at the three aspects that lay the foundation of P2P computing technology. This will be followed by an attempt at formulating definitions for some of the concepts associated with this technology. Expanding on these core concepts, a closer look will then be taken at the application of P2P computing in the business and tertiary-education environments. Following, a historical overview of the P2P computing phenomenon.

4.2 P2P: A historical perspective

As indicated in section 4.1, P2P computing technology is not a new concept. P2P in its basic form had, in fact, come into being when the Internet started taking shape in the 1970s. Today's P2P computing infrastructure is, therefore, not merely the culmination of everything accomplished in the 1970s but also the sum of the evolutions of these earlier experimental phases. In order fully to grasp and appreciate today's P2P technology, a closer inspection is needed of the historical events leading up to and culminating in modern-day P2P applications. According to Barkai (2002: 39), the historical foundation on which modern-day P2P applications rest can be divided into the following three components:

- Direct exchange
- Distributed processing
- Online collaboration.

It should be noted, that even though these three components have evolved independently from each other, the sum total of their combined outputs created the powerful computing environment that we know today. In the following sections, each of these three components forming the historical foundation of P2P will be discussed briefly.

4.2.1 Direct exchange

As discussed in chapter 2, the Advanced Research Projects Agency Network ("ARPANET", for short) was the predecessor of the modern-day Internet. Some of the earliest and most important research developments that influenced the present-day distributed-computing environments were effected at Xerox's Palo Alto Research Centre ("PARC", for short) in the 1970s. These research efforts, in turn, gave rise to the single-user workstation, file and print servers, Remote Procedure Call ("RPC", for short) and high-speed resource-sharing local networks (Barkai, 2002: 40).

As was pointed out in chapter 2, the Internet was originally designed with direct exchange in mind. The emergence and development of a messenger service were the logical next steps in the evolutionary path of the Internet. The first real-time communication over the Internet took place in August 1988. It was not P2P exactly, though, as a server was used to host the conversations. What it did give rise to, however, was the Internet Relay Chat facility ("IRC", for short), which was created due to the need to send real-time messages to other Internet users, which makes it the most immediate forerunner of today's instant-messaging systems (Miller, 2001: 210).

4.2.2 Distributed processing

As soon as computers are interconnected and communicating, we have a "distributed system". This approach to computing has been known by several names, such as "metacomputing", "scalable computing", "global computing", "Internet computing" and, lately, as "Grid computing" (De Roure, 2003: 65-66).

According to Metz (2002), a network of computers collaborating on a common task, each donating processing power, storage space and/or other resources, can be described as a "grid".

Furthermore, grid computing is a hard- and software infrastructure that clusters and integrates high-end computers, networks, databases and scientific instruments from multiple sources to form a virtual supercomputer on which users can work collaboratively (Lais, 2002).

In essence, the term "grid computing" has been coined to denote a proposed computing infrastructure for advanced science and engineering (Foster, 2003: 171).

A cousin to both cluster computing and parallel processing, grid computing can be thought of as distributed and large-scale cluster computing, or as a form of network-distributed parallel processing (Thilmany, 2003).

For an aggregation of computers to be a grid, therefore, it must do the following (Lais, 2002):

- Co-ordinate resources that are not subject to centralised control.
- Use standard, open, general-purpose protocols and interfaces.
- Render quality service.

Next, a look at the third historical basis upon which P2P applications are built.

4.2.3 Online collaboration

As was indicated in section 4.2.1, online collaboration existed long before it was labelled "P2P computing". As was pointed out, Internet Relay Chat ("IRC", for short) and Bulletin Boards ("BB", for short) facilities provided the first mechanisms for one-to-one direct exchange of information. The IRC facility provided a mechanism for individual users to communicate and share ideas, and files in an instantaneous fashion through media such as text, video and voice.

The previous sections were used to focus on the historical foundations of P2P computing, thereby laying a foundation for a more detailed analysis of the concept of P2P computing.

4.3 P2P: A modern-day perspective

If a million people use a website simultaneously, doesn't that mean that we must have a heavy-duty remote server to keep them all happy? No; we could move the site onto a million desktops and use the Internet for coordination (Gelernter, 2000).

4.3.1 Defining "P2P"

Depending on the way in which the term is defined, P2P can even be seen as one of the oldest architectures in the world of telecommunication. The telephone system, the discussion forums of Usenet or the early form of the Internet can, in fact, all be classified as P2P systems (Schoder, 2003: 27). For this reason, before embarking on a detailed analysis of P2P, the phrase "Peer-to-Peer" needs to be defined first.

4.3.1.1 A simple definition

What is "P2P"? Is P2P a set of protocols, an IT architecture, a design philosophy stressing decentralisation, a business model or merely a fad?

In reply to the above question, a simple definition is offered: "It (P2P) refers to the topology and architecture of the computers in a system" (Bricklin, 2000), "in which (system) each computer has equivalent capabilities and responsibilities" (Miller, 2001: 19). As is evident, the foregoing definition is in glaring contrast to the traditional client/server network architecture, "for each participating computer or node in a P2P system is called a 'peer', meaning that the participants interact as equals" (Kubiatowicsz, 2003: 33).

The foregoing definition will not suffice, however, as P2P encompasses far more than what has been mentioned thus far.

4.3.1.2 A complex definition

The best way to characterise P2P systems is not by a common technology, but by what it does (Lee, 2002). With P2P, computers share data and resources, such as spare computing cycles and storage capacity, via the Internet or

private networks. Machines can also communicate directly and can perform computing tasks without using central servers (Clark, 2001: 18).
With the advent of P2P computing, the distinction between server, client and router has become blurred, because individual computers sustain these roles, communicating and sharing resources without dedicated servers (Lee, 2003: 49).

Expanding on the foregoing descriptions, a P2P system can be identified by the following five salient features (Mauthe, 2001; Miller, 2001: 19):

- Thanks to their structure and organisation, P2P systems are inherently fault tolerant.
- P2P systems are self-organising in the sense that the various system components work together without any central-management instance assigning roles and tasks.
- The network facilitates real-time transmission of data or messages between the peers.
- Peers can function as both client and server.
- The peers provide the primary content of the network.
- The network gives control and autonomy to the peers.
- The network accommodates peers that are not always connected and that may not have permanent Internet Protocol ("IP", for short) addresses.

Subjecting a specific service or application to the above five criteria can help to identify the relative P2P activity. In addition, Miller (2001: 22) subjected these criteria to a list of Web services to determine whether or not it satisfied the definition of P2P computing. The Web services included the following: Distributed computing, File-sharing/-swapping, Group collaboration/ Conferencing, Instant messaging, Internet telephony, e-mail, File Transfer Protocol (FTP), Internet chat, Usenet newsgroups and Web browsing. Web services can be regarded as a complementary concept that will enrich the Internet infrastructure and its utilising applications, including any P2P application (Schoder, 2003: 27).

The results are summarised in table 4.1 on the next page.

Table 4.1: Criteria for identifying P2P activity (Miller, 2001: 22)

Criterion	Distributed computing	File-sharing/ -swapping	Group collaboration/ Conferencing	Instant messaging	Internet telephony	E-mail	FTP	Internet chat	Usenet newsgroups	Web browsing
Enables two or more peers to communicate with each other in real time	Yes	Yes	Yes	Yes	Yes	No	No	Yes	No	No
Peers function as both client and server	Yes	Yes	Yes	Yes	Yes	No	No	Yes	No	No
Content provided/hosted by individual peers	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes	No	No
Gives control and autonomy to individual peers	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No
Recognises variable connectivity and temporary addresses for both computers	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No
Peer-to-Peer?	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No

As indicated in the above table, some Web services, including e-mail, FTP, Internet chat and Usenet newsgroups, which can be classified as services related to the World Wide Web, cannot be considered as P2P computing.

4.3.1.3 P2P computing *vs* P2P network *vs* P2P communications

Related terms that are sometimes incorrectly used as though interchangeable with the term "Peer-to-Peer computing" are "Peer-to-Peer networking" and "Peer-to-Peer communications" (Niemand, 2003). A *P2P network* allows every computer in the network to act as a server to every other computer in the network. Dedicated servers may or may not be present (Barkai, 2002: 19).

P2P communications, in turn, enables both participants to initiate, manage and terminate the session.

It would, therefore, be safe to say that a P2P network implies the P2P communication between the computers in the network. It should be noted, however, that the P2P communication could also occur between computers in a network that is not a P2P network.

In terms of *P2P computing*, not every computer has to be a server to every other computer. The relationship between the users in P2P computing is arranged in some kind of manner. Computers in P2P computing have software layers that provide server services. The peers in P2P computing are likely to exercise P2P communications; that is, they have equal authority and responsibility regarding communications with each other. The peers may, however, also agree on communication arrangements that are not symmetrical. For P2P computing, communications is merely a means to facilitate sharing of resources.

In conclusion, the phrases "P2P network" and "P2P communications" refer to the physical network and the mode of communication. In contrast with these, the phrase "P2P computing" refers to an end-user application-level environment in which the typical file-sharing takes place.

4.4 Objectives of P2P technology

As with any computing system, the objective of P2P systems is to support applications that satisfy users' needs. Selecting a P2P approach is often driven by one or more of the following objectives (Milojicic, 2002):

- Cost-sharing/Cost reduction
- Anonymity/Privacy
- Improved scalability/reliability
- Dynamism
- Resource aggregation and interoperability
- Increased autonomy
- Enabling *ad hoc* communication and collaboration.

Following, a closer look at each of the foregoing objectives.

4.4.1 Cost-sharing/Cost reduction

In a centralised system, that is, in the typical client/server architecture, serving all the clients can be seen as the largest contributing factor to the cost incurred in running such system. In a P2P architecture, on the other hand, the cost can be spread across all the peers. In addition, most of the cost-sharing is realised by the utilisation and aggregation of unused resources, thus increasing productivity even more and empowering the workers themselves to do more (Fattah, 2002: 12).

4.4.2 Improved scalability/reliability

Given the lack of central authority for autonomous peers, improving system scalability and reliability is one of the main objectives. As a result, algorithmic innovation in the area of resource discovery and search has been a clear area of research, giving rise to new algorithms for existing systems and the development of new P2P platforms. In conclusion, as summed up by Kubiatowicsz (2003: 34), the system needs to work well "with thousands, millions or even billions of clients".

4.4.3 Resource aggregation and interoperability

The greatest potential of P2P computing lies in its ability to yield the computing power of one's network to one's employees, one's partners and even one's customers (Fattah, 2002: 12). In contrast with the traditional client/server architecture, the P2P computing model can be seen as a more decentralised approach. The natural tendency within this decentralised architecture is for resources naturally to aggregate.

Large-scale projects that harness many thousands of computers jointly to perform a Herculean computational task are, for instance, SETI@Home, distributed.net and Endeavours. This activity is referred to as "cycle-sharing" (Barkai, 2002: 229). The aggregation occurs when each computer in the P2P system adds additional computing power and/or storage space to the P2P system. File-sharing systems, such as Napster and Gnutella, are online aggregators of resources. In such cases, it is both disk space to store the community's collection of data and bandwidth to move the data that has been aggregated (Milojicic, 2002).

4.4.4 Increased autonomy

In many cases, users of a distributed system are unwilling to rely on any centralised service provider. Instead, they prefer all data collected and work done on their behalf to be performed locally. P2P systems support this level of autonomy, simply because they require the local node to do work on behalf of its user.

Prime examples of this are the various file-sharing systems, such as Napster, Gnutella and FreeNet. In each case, users are able to obtain files that would not be available at any central server because of licensing restrictions. Individuals autonomously running their own servers, however, have been able to share the files because they are more difficult to find than a server operator would be.

4.4.5 Anonymity/Privacy

With any system, a user or client may not want anyone or the service providers to know about his/her involvement and participation in the system. In a typical client/server architecture, this level of anonymity is not always possible, for the server will typically be able to identify a user, at least by his/her related Internet Protocol ("IP", for short) address. One of the grand ideas behind P2P systems, however, is their potential to provide a "cloak" under which users could conduct information exchanges without revealing their identities or even the nature of the information they are exchanging (Kan, 2001: 117).

4.4.6 Dynamism

Dynamic participation has both a philosophical and a practical advantage. The loosely connected structure and explosive growth of the Internet suggest that any P2P system must be similarly flexible and dynamic in order to be scalable and to sustain long-term use (Dingledine, 2001: 275). The system of computers or nodes must, therefore, be able to "enter" or "exit" the network on a continuous basis.

4.4.7 Enabling *ad hoc* communication and collaboration

Owing to the fact that a P2P system does not typically rely on an established infrastructure (that is, a server), it is ideally designed to support a dynamic and *ad hoc* environment. The users/clients can enter or exit the system at will, according to their changing needs, thereby enabling communication and collaboration within the P2P community.

4.5 P2P architecture

Before analysing the architecture of a P2P system, it should be noted that all computer systems can be classified into either a distributed or a centralised system (Milojicic, 2002), as depicted in figure 4.1 on the next page.

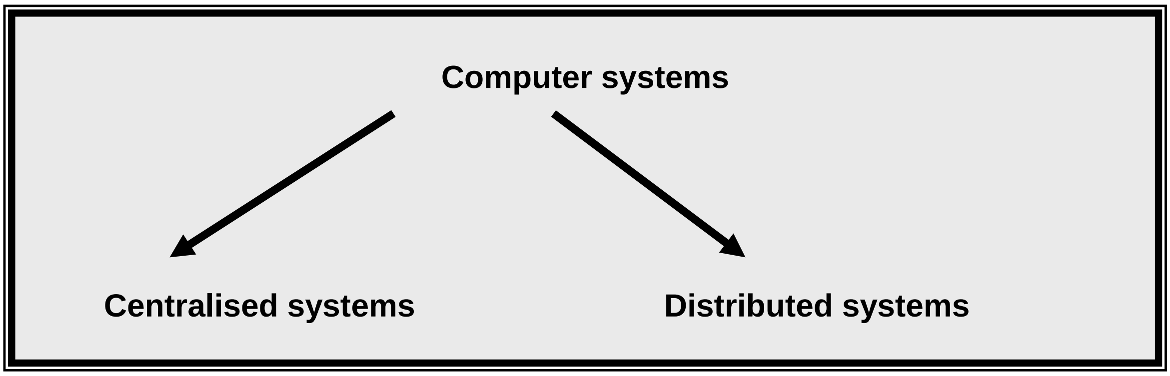

Figure 4.1: A taxonomy of computer systems (Milojicic, 2002)

Within this taxonomy of computer systems, the distributed systems can be sub-divided into the Client/Server model and the P2P model, as in figure 4.2 below:

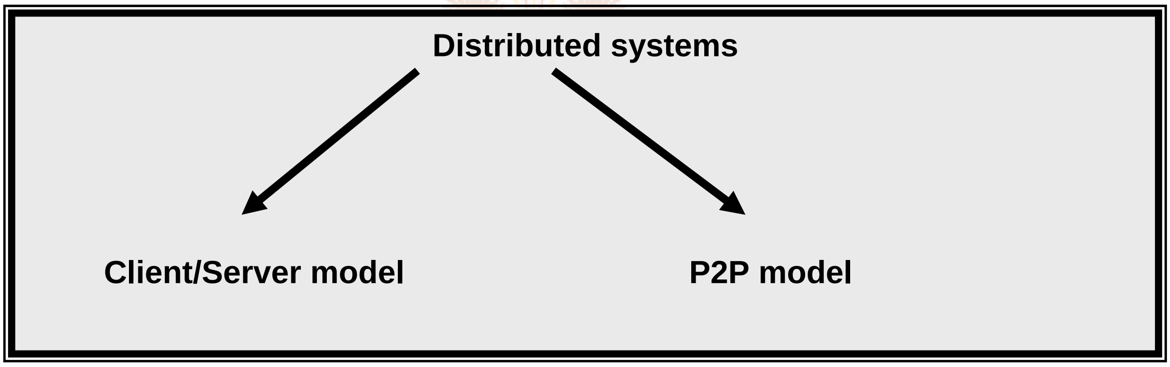

Figure 4.2: A taxonomy of the distributed systems (Milojicic, 2002)

The following section will be devoted to taking a closer look at some of the P2P-architectured models.

4.6 P2P models

One of the great strengths of P2P-based models is their decreased dependency on the server and their decentralisation of control from servers. In the conventional client/server architecture, however, the client is subject to various rules and regulations, but in terms of the P2P model environment, the user/client enjoys greater control. According to the Dreamtech Software Team (2002: 2-5), P2P models can be divided into the following categories:

- Pure P2P
- P2P with a simple discovery server
- P2P with discovery and lookup servers
- P2P with discovery, lookup and content servers.

Following, a brief discussion on each of the foregoing categories.

4.6.1 Pure P2P

The pure P2P model consists of and is entirely dependent on personal computers. Anon (2001) concurs and adds that it is the "…truest P2P architecture". This revolutionary architecture is contradictory to the "old" networking model, which consisted of personal computers and servers. Typically, these categories of networks do not use servers to effect routing, caching, file-sharing, recording of search results or any other peer activities (Moore, 2002: 107) and must, therefore, be effected by the peers themselves (Miller, 2001: 379). After having downloaded the relevant P2P application onto the memory of the machine, a dynamic search for other peers in the network of P2P application is undertaken and a connection is established. In this way, the entire communication process, that is, the

- transferring of data in the form of uploading and downloading files
- carrying out of online activities
- sending of requests
- receiving of responses, etc,

occurs without any assistance from servers.

In essence, the pure P2P model allows users/clients to create and determine their own unique networking environment.

4.6.2 P2P with a simple discovery server

According to Anon (2001), this may also be deemed a "user-centred application". As the name suggests, a server functions as a basic administrator. The primary role of this server is to provide the names of peers already connected to the incoming peer, who notify the server of their existence by logging on. It should be noted here, however, that the server only provides the names of connected peers, whilst the establishment of connection and the actual communication remain the responsibility of the peers. By using a discovery server, a peer eliminates the problem of finding that first peer to connect to (Miller, 2001: 380).

For resource downloading, a peer has to approach each individual peer in this specified environment until the resource is found. This is in contrast with the original client/server architecture, where all the relevant content was located on the server, thereby rendering the search for resources on the P2P with a simple discovery server, as with pure P2P, a relatively time-consuming process.

4.6.3 P2P with discovery and lookup servers

In terms of this model, the features of the pure P2P and P2P with simple discovery server are enhanced with another feature of a server, that is, that of looking at the resources available in the P2P network. Anon (2001) labels this as a "…data-centred approach". This architecture or approach is an enhancement compared to the P2P with basic discovery server, for it reduces the amount of time wasted by visiting each individual peer to obtain resources and to glean information from each peer.

A lookup server contains a central index of all the files stored on all the peers in the network. When one peer submits a data request, the central index is queried, the hosting peer is found and the requesting peer is connected to the peer hosting that file. For the lookup server to function, each peer must upload a list of its file contents to the central index. This is typically done when the peer first logs onto the network, as part of the normal connection process (Miller, 2001: 381).

4.6.4 P2P with discovery, lookup and content servers

As in the typical client/server architecture, all the facets of requesting and acquiring information and resources resided with the server, hence the convergence of the above models with Web architectures and infrastructure (Anon, 2001). In addition, the individual peers are not allowed directly to connect to each other. Should a peer require a specific resource, it would communicate with the server, whereupon the server would, in turn, process the request and display the information.

A possible drawback of the said architecture scheme could be slower processing of requests by the server as the volume of requests increases. Higher costs and the increased possibility of single-point failure (as the server is responsible for handling all processes) can be identified as further possible drawbacks of this kind of system. This type of architecture closely resembles a traditional client/server network (Miller, 2001: 381).

The models described above serve to illustrate three of the four main types of models that have evolved in the P2P environment. All three models share certain features – a user may act either as a server (hosting files for others to download) or as a client (downloading from hosting servers). Each model creates an environment that is conducive to the ready dissemination of information.

The difference between the models lies in the manner in which each framework has been laid out. In the centralised framework, all users surround a central server, thus creating a framework with a "star-like" quality, with the nucleus (the central server) having various rays (the users) radiating from it.

In the decentralised framework, users and servers are all connected to each other, thus creating a "vine-like" shape for the layout of its framework. In a controlled decentralised framework, the framework shape is a unique blend, giving it a "web-like" quality.

All P2P models share the same objective of sharing information within their frameworks. How the information is shared and how much of it is shared are some of the main differences between each P2P model (Software & Information Industry Association, 2001). The next section will be devoted to a closer look at applications utilising P2P computing.

4.7 Applications utilising P2P computing

While the previous section was used to introduce the reader to the basic framework and architecture enabling modern-day P2P computing, this section will, in turn, be devoted to analysing some of the existing applications utilising P2P computing.

When mentioning the word "P2P" in conversation, most people will respond with saying "Napster". Even though Napster probably is the best-know type of P2P application, it is by no means the only or the largest P2P application. There are, indeed, at least five distinct types of P2P applications in use today. Shirky in Oram (2001: 22) concurs with this and adds that if the answers to both the following two questions were "yes", it would be a P2P application, but if the answers were “no”, it would not be a P2P application:

- Does it allow for variable connectivity and temporary network addresses?
- Does it give significant autonomy to the nodes at the edge of the network?

File-sharing (the category that Napster falls into) is just one type of application; the other types include instant messaging, distributed search, group collaboration and distributed computing, as depicted in figure 4.3 (Miller, 2001: 30). Next, a discussion on each of these applications.

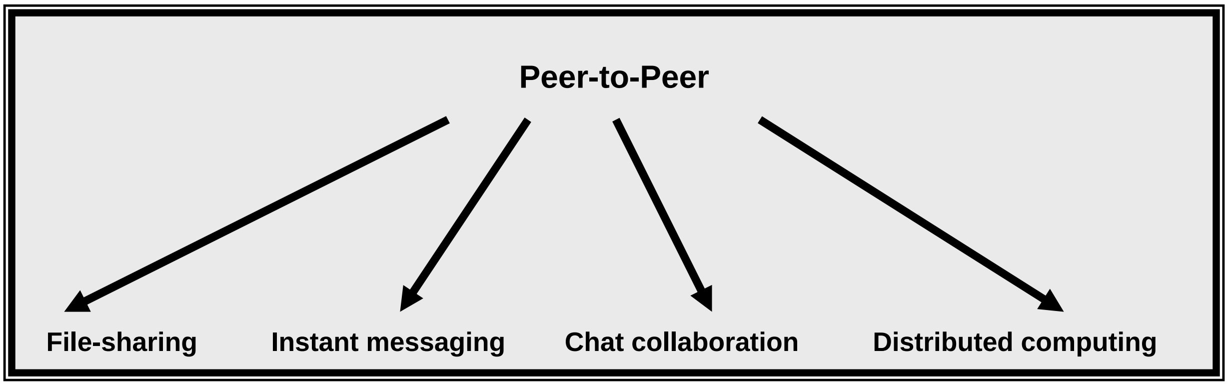

Figure 4.3: P2P application types (Miller, 2001: 30)

4.7.1 P2P communications (instant messaging)

At more than 200 million users, instant messaging is by far the most popular current P2P application – even when most instant-messaging users are unaware of the fact that they are, indeed, engaging in P2P computing (Miller, 2001: 30).

Instant messaging differs from the ordinary e-mail facility in the immediacy of the message exchange. In addition, most exchanges are text-only. Some services, however, such as America Online (“AOL”, for short), allow for voice messaging and file-sharing (WhatIs.com, 2004). Instant messaging can, therefore, be seen as a pure form of P2P computing.

Inside the instant-message network, the messages flow from one user to another without interference by or help from any servers whatsoever. It should also be noted that the instant-messaging environments operate on their own proprietary instant-message network.

According to Miller (2001: 31-32), the following can be deemed key players in the instant-messaging environment:

- American Online (“AOL”, for short) Instant Messenger (www.aim.com)

- ICQ (web.icq.com)
- Microsoft (MSN) Messenger (messenger.msn.com)
- Yahoo! Messenger (messenger.yahoo.com).

4.7.2 P2P file-sharing

As noted in the previous section, most people invariably think of Napster and MP3 file-downloading upon hearing the word "P2P". The truth is, however, that P2P file-sharing entails much more than Napster or MP3 music files.
A file could contain any type of information, such as a school history report, a business case or a travel journal. Files could also contain multiple formats; for example, a text document could include embedded graphics (Moore, 2002: 37).

The traditional definition of "file-sharing" is that "two or more users can access (read and/or modify) the same file, typically the same single instance of a file". This is a capability that a typical shared distributed-file system supports. File-sharing in the P2P sense has come to mean "the transfer of a copy of the file" (Barkai, 2002: 299-300). The success of file-sharing depends on both distributed-content storage and dynamic search (Quintana, 2001).

Next, a brief discussion on the basics of P2P file-sharing.

4.7.2.1 P2P file-sharing in a nutshell

Two computers (A and B) connected to each other can copy files from one computer (A) to the other computer (B) on condition that the two computers in question be connected in a network. Computer A, that is, the computer with the file on a specific drive, needs to be configured so that the specific folder containing the file could be shared. In its turn, computer B must be configured to ensure that it knows which computer holds the specific folder containing the required file. Computer B then accesses computer A and transfers the file to its own drive.

Since the Internet is nothing more than a giant network, namely the "Global Area Network" (or "GAN", for short), there is no reason why any two computers connected to the Internet could not share files, just as two computers on a Local Area Network ("LAN", for short) can. In terms of the P2P system, an additional software layer enables the computers at the edge of the network to act as clients and/or servers, as circumstances may demand (Quintana, 2001).

This process of connecting two computers over the Internet and of then copying files between them is what P2P file-sharing is all about (Miller, 2001: 97).

Something to note in this regard is the fact that a P2P file-sharing network operates outside the traditional Internet Domain Name Server ("DNS", for short) addressing system. This enables each peer on the network to connect irrespective of which Internet Service Provider ("ISP", for short) is used and, in some cases, this ensures anonymity for the individual personal computer. After the connection to the specific P2P network has been established, the next step in file-sharing is taken, namely the search for a specific file. The P2P model will determine how this objective will be reached.

A pure P2P model does not facilitate a central-index server, which is why finding files is embarked upon in real time. In contrast with the pure P2P model, the "brokered" models utilise the creation of an index of useable files on servers. In this way, the search for a specific file is only undertaken on the server.

Once the specific file has been located, it needs to be copied to the requester's PC. The transfer is effected via the Hyper-Text Transfer Protocol ("HTTP", for short). In fact, during the transfer of the file, the computer with the copy of the file on it is turned into a temporary HTTP server. Each bite and byte of the file is sent to the requesting PC via the Internet. Upon having downloaded successfully, access to the temporary HTTP server (or computer) is terminated.

Following, a list of some of the key players in the P2P file-sharing environment (Mitchell, 2004):

- eDonkey/Overnet
- Shareaza
- WinMX
- eMule
- BitTorrent
- Kazaa.

4.7.2.2 Issues surrounding P2P file-sharing

Due to the dynamic nature of P2P computing, new developments are being effected in this study field almost on a daily basis. The major issues surrounding the P2P field of study include the following (Von Lohmann, 2003):

- Issues with regard to copyright
- Ethical issues.

(a) Issues with regard to copyright

The future of P2P file-sharing and related technologies is inextricably linked with that of copyright law (Von Lohmann, 2003). While the concept of "copyright law" may sound simple, in reality the court cases surrounding P2P-sharing (with specific reference to Napster) prove quite the opposite. This complexity can be attributed to the fact that the Internet and P2P applications are used in a global manner; in some cases, P2P applications offer the majority of users anonymity. In addition, copyright laws vary from one country to the next, thus allowing P2P providers easily to set up in whatever country they would not be breaking the law (Burke, 2003).

According to the *Statutes* of the Republic of South Africa, the following obtains to copyright:

An "author" in relation to

- a computer program is "the person who exercises control over the making of the computer program".

A "computer program" means

- "a set of instructions fixed or stored in any manner and which, when used directly in a computer, directs its operation to bring about a result".

A "copy" means

- "a reproduction of a work, and, in the case of a literary, musical or artistic work, a cinematography film or a computer program, also an adaptation thereof; provided that an object shall not be taken to be a copy of a work of architecture unless the object is a building or a model of a building".

According to the *Statutes*, copyright shall obtain to the following bodies of work, should they be original:

- Literary works
- Musical works
- Artistic works
- Cinematography films
- Sound recordings
- Broadcasts
- Programme-carrying signals
- Published editions
- Computer programs.

The concept of "copyright" may be defined as follows (UCT, 2000):
"Copyright is part of a group of intellectual-property rights, which (rights) provide legal protection to creators of works of the mind."

In South Africa, copyright is governed by the Copyright Act No 98 of 1978, as well as by the Regulations made in terms thereof, and it grants owners of copyright (authors and other creators of intellectual property) the right to

- reproduce the work
- create derivative works based on the original work
- distribute copies of the work
- perform the work
- display the work in public.

In literary artistic works other than photographs, copyright endures for the lifetime of the author and for 50 years after his/her death, or for 50 years after the posthumous first publication of his/her work. In a film, a photograph or a computer program, copyright endures for 50 years from the end of the year in which the work is made available to the public or after the work is first published, whichever is the longer period.

The author(s) or holder(s) of his/her/their copyright can take legal action where there is an infringement of his/her/their rights. The remedies provided include delivery of the infringing material, damages and an interdict preventing further infringement of his/her/their rights. The courts have the power to award additional damages where there has been a flagrant infringement of copyright. The Copyright Act also makes provision for criminal penalties – a fine (of R5 000) and/or imprisonment of up to three years per infringement for a first conviction. The maximum fine and/or imprisonment penalty for a second conviction is R10 000 and/or five years per infringement (UCT, 2000).

Copyright and Peer-to-Peer

The nature of digital file-sharing technology inevitably impacts on copyright laws. As the Napster saga illustrates, the future of P2P file-sharing is inextricably linked to copyright law (Von Lohmann, 2001). In the legal arena, copyright owners are not only targeting the makers of file-sharing software, such as Kazaa and Napster, but they are also the companies providing products and/or rendering value-adding services, such as MP3Board.com. The lesson (for any P2P developer) learnt from these early legal battles is to ensure that a solid legal strategy be put in place before commencing the development of the P2P application or add-on. In real-world terms, software developers cannot always control or predict the behaviour of their end-users. Closely related to the copyright issue is the issue of the ethical use of P2P systems. Following, a brief discussion on the ethical issues surrounding P2P computing.

(b) Ethical issues

Given all the media hype generated by the Napster court cases, individuals and organisations have placed new emphasis on the ethical use of information technologies. According to Hellriegel (1996: 631), the reason for this is that the ethical issues surrounding computers arise from their unique technological characteristics:

- Computers make mistakes that no human being would make.
- Computers communicate at high speed over great distances and at low cost.
- Computers have huge capacities quickly and economically to store, copy, erase, retrieve, transmit and manipulate information.
- Computers have the effect of radically distancing (depersonalising) originators, users and subjects of programs and data from each other.
- Computers can collect and store data for one purpose, but the selfsame data can just as easily be used for another purpose. Data can also be stored almost indefinitely.

"Ethics" can loosely be defined as "a set of rules or values that defines a community's sense of right and wrong". This definition corresponds with that formulated by Hellriegel (1996: 17), in terms of which ethics is defined as "a set (or code) of moral principles, values and conduct that decision-makers apply to issues that are not specifically addressed by law". In the context of P2P computing ethics, ethics can be considered a double-edged sword, that is, the ethical use of P2P technology by the end-users and the ethical use of P2P technology by the developers. The rise in interest in computer ethics saw the formation of The Computer Ethics Institute, where the following "Ten Commandments" of computer ethics were formulated.

Table 4.2: The "Ten Commandments" of computer ethics (Hellriegel, 1996: 631)

1.	**Thou shalt not use a computer to harm other people.**
2.	**Thou shalt not interfere with other people's computer work.**
3.	**Thou shalt not snoop around in other people's computer files.**
4.	**Thou shalt not use a computer to steal.**
5.	**Thou shalt not use a computer to bear false witness.**
6.	**Thou shalt not copy or use proprietary software for which you have not paid.**
7.	**Thou shalt not use other people's computer resources without their authorisation or without compensating them properly.**
8.	**Thou shalt not use other people's intellectual outputs.**
9.	**Thou shalt think about the social consequences of the program you are writing or the system you are designing.**
10.	**Thou shalt always use a computer in ways that demonstrate consideration and respect for your fellow human beings.**

P2P networking allows one's network to be open to various forms of attack, break-in, espionage and malicious mischief. P2P, however, does not bring any novel threats to the network, just familiar threats such as worms and virus attacks. P2P networks can also allow an employee to download and use copyrighted material in a way that violates intellectual-property laws, and to share files in a manner that violates the security policies of an organisation.

Applications such as Napster, Kazaa, Grokster and others have been popular with music-loving Internet users for several years now, and many users take advantage of their employers' high-speed connections to download files at work. This presents numerous problems for the corporate network, such as using expensive bandwidth and being subject to a virus attack via an infected file download.

Type of content shared in P2P file-sharing

The following are a few examples of the file types commonly shared on Kazaa and many other P2P applications:

- Audio, such as, MP3 (music)
- Video files, such as, MPEG (cinematography)
- Images, such as, pictures (JPEG)
- Documents such as books
- Software such as MS Office.

Active sharing is the cornerstone of a useful P2P experience. In order for users to benefit from such collaboration, however, they all need to share files in accordance with the end-user licensing agreement. Since successful P2P computing is a two-way street of file up- and downloading, it is prone to abuse by "freeloaders".

A "freeloader" can be defined as "an individual peer in the P2P computing environment (who) downloads more resources than he/she uploads in the system". In other words, a freeloader habitually reaps the rewards of other peers' collaboration. It should also be noted here that these freeloaders or "Net leeches" defeat the very objective of a fully decentralised system and must, therefore, be deemed a potential cause of the downfall of P2P systems (Lee, 2002).

P2P systems may, however, combat the freeloading effect by incorporating coercive measures into their design and deployment, thus ensuring that users donate resources (Dingledine, 2001: 292). This sentiment is shared by Lee (2002), who goes further to suggest that these systems be equipped with anti-freeloading mechanisms such as tying the number of downloads to the number of uploads, or using a point system actually to reward all contributors.

4.7.3 P2P group collaboration

Modern-day P2P group collaboration has its roots in yesteryear's groupware.

For this reason, it seems appropriate to define “groupware” as “a general term (that) denotes software-based tools which can be used to support a distributed set of workers”. This covers applications as incongruent as Windows for Workgroups and PC videophones (Norris, 1996: 279).

In the client/server architecture, the server was responsible for the hosting of all the groupware applications and data. The server was also responsible for group scheduling and communications. The benefits of implementing a P2P group-collaboration application freed up valuable server space and processing power, ultimately deriving a huge cost benefit.

Since P2P group collaboration targets a market that readily pays for new technology, that is, the corporate environment, this section of P2P applications can be deemed the prime field for future P2P development and growth.
The basic functioning of P2P collaboration can best be explained by comparing it to the more traditional client/server collaboration (please see table 4.3 below):

Table 4.3: Comparison between client/server and P2P collaboration (Milojicic, 2002)

Client/Server	Peer-to-Peer
Does not allow direct user-to-user transactions.	Allows direct user-to-user transactions.
All communication, scheduling and data-storing is done and controlled by a central server.	All communication, scheduling, file- and application-sharing is done on a P2P basis.
Limited, to a certain extent, to the boundaries of the organisation.	Communication and transfer can be effected on an internal network and on the Internet.

P2P groupware combines features from many different P2P (and non-P2P) applications. The ideal P2P groupware system would be part Napster, part AIM and part Outlook – and would be accessible from anywhere in the world (Miller, 2001: 273).

The ideal P2P groupware system should, in essence, contain the following:

- E-mail
- Instant messaging
- Message boards
- Real-time chat rooms
- Whiteboarding
- Calendar

- Scheduler
- Contact management
- File-sharing
- Application-sharing
- Application development
- Database management
- Project management.

It should be noted here that, for all of the above applications in a P2P collaborative system, success will largely depend on the management of the content in the system. Because the individual members in the system will supply all the data and information, the content will be unstructured in nature. The P2P collaborative system will have to reorder and restructure the content provided.

According to Miller (2001: 275-276), there are only two big players in collaborative groupware, namely

- Lotus (IBM), and
- Microsoft.

4.7.4 P2P distributed computing

The last type of P2P computing (distributed computing) is in many respects very different from the applications discussed thus far. The latter types of P2P computing always involved one personal computer connecting to another personal computer for messaging, file-sharing and collaboration.

Even though 50 years of innovation have increased the raw speed of individual computers by a factor of around one million, they are still far too slow for many challenging scientific problems (Foster, 2000). Before P2P distributed computing, highly complex computing problems were solved by sheer force. Large “supercomputers” were set aside to tackle the problems by crunching calculation after calculation until a solution was found.

With P2P distributed computing, however, the highly complex computing problem gets broken down into various smaller calculations, which calculations are then fed into a network of individual PCs, each responsible for processing its calculation independently. Upon finishing their individual processes of calculation, the results are fed back to the central server, which combines the results obtained from the individual computers to formulate the result for the project.

It should be noted, too, that the above distributed-computing process could be affected either internally or externally, with “internally” being defined as “within the boundaries of a single business, across a closed network of computers and with absolute control over the network of computers”.

In contrast with the definition of “internally”, the term “externally” refers to using the Internet to connect computers in a global fashion and to effecting most of the processing when users’ computers are offline. In this way, the specific project must be downloaded from the project Website. The processing occurs when the individual computer is idle. Once the computation has been completed, the results are reported back to the project Website.

According to Miller (2001: 339), some of the potential users of distributed computing include

- energy companies undertaking geographic analyses before drilling for gas or oil
- manufacturers performing complex structural analyses
- engineering firms embarking on stress-tests for large projects
- entertainment and graphics companies rendering computer-generated animation
- biotechnology firms modelling various viruses
- pharmaceutical companies developing new drugs and vaccines
- generic research (think of efforts such as the Human Genome Project)
- any company, university or independent project involving complex mathematical modelling.

It is evident from the above examples that some projects are of a charitable nature, whilst others are purely commercial.

4.8 Benefits of Peer-to-Peer computing

Following, some of the benefits that could be derived from P2P computing (Jiangang, 2004; Cugola, 2001; Miller, 2001: 27-28):

- A peer does not need to deposit an object that it wants to offer in one or more servers beforehand.

- A peer who needs the object directly requests the provider peer to provide the consumer peer with the object.
- P2P ensures that information and services are no longer gathered at a single point of accumulation.
- P2P leverages previously unused resources found on the hundreds of millions of computers (and other devices) that are connected to the edges of the Internet.
- P2P enables faster delivery of information from one computer to another by bypassing a central server.
- P2P frees up bandwidth on the Internet (or on a private network). Under the traditional client/server model, the server is the bottleneck and often cannot deal with everything at once.
- P2P increases personal efficiency and personal empowerment. Users no longer have to wait in queues to perform essential tasks, since all activities take place at the user's discretion.
- P2P represents a significant cost saving over the client/server models. Since resources and computing power is distributed across the entire network, there is no need for expensive centralised severs. This also reduces the need for centralised management, storage and other related resources.
- P2P requires no centralised managing, overseeing or controlling.
- P2P offers easy scalability, so that all that is needed for a network to grow is to add more peers.
- P2P increases the fault tolerance of a network. Since no one part of the system is essential to its operation, one could take down a few nodes and the network would still remain functional.
- P2P offers increased privacy, since all data and messages are directly exchanged between two computers.
- P2P in networks is more flexible and adaptable than traditional client/server networks.

Even though the benefits to be derived from P2P are numerous, it should be noted that some disadvantages could also be identified for P2P computing. The next section will, therefore, be devoted to identifying some of the possible drawbacks of P2P.

4.9 Drawbacks of Peer-to-Peer computing

Although the benefits to be derived from P2P computing are numerous, the fact remains that the P2P model cannot possibly be considered the best option or solution for every scenario. The following, then, would be instances in which certain drawbacks of Peer-to-Peer would come into play (Jiangang, 2004; Miller, 2001: 28):

- A complex protocol is needed to maintain a logical P2P network.
- In a P2P network there is no guarantee that content/resources will always be available – a peer could go offline if someone were to shut off his/her computer.
- A P2P network has no way of enforcing content ownership (copyright).
- A P2P network is unable to enforce standards (either technological or ethical/moral/social).
- In a P2P network, it is extremely easy to propagate all sorts of undesirable items and activities, including misinformation.
- The P2P model increases the chances for a network (or an individual system, at that) to become exposed to hackers, attacks, viruses and other malicious damage.
- An unprepared P2P network could be overwhelmed by increased traffic.
- P2P is dogged by a lack of standards, infrastructure and support.
- P2P transactions are difficult to translate into revenue streams – and this lack of revenue generation could hinder future development.

By identifying the possible drawbacks in P2P computing it is possible to embark on an analysis of some of the potential obstacles in this technology.

4.10 P2P environments

According to Milojicic (2002), the target environments for P2P consist of the Internet, intranets and *ad hoc* networks. P2P systems connected via the Internet support connections in the spectrum from dialup lines to broadband. The underlying architecture can rely on personal home computers, corporate desktops or personal mobile computers (notebook and handheld computers).

There are three main markets for Peer-to-Peer:

- Consumer space encompasses the use of P2P for personal use, such as music and content-sharing, instant messaging, e-mail and games.
- Enterprise space or business P2P applications include biotech, financial, traditional IT and Business-to-Business solutions. Examples of 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.

For the purposes of this study, however, the focus will fall on the implementation of P2P technologies in businesses and tertiary-education institutions.

4.10.1 Implementation in tertiary-education institutions

As indicated in this study, the fundamental Internet architecture has and continues to revolutionise modern society. Fast-developing computer technologies and communications offer a means of creating new learning activities and experiences (Kook, 2001: 39). According to Guscott (2001), we should make use of these new P2P technologies to help fulfil our mission, which is to provide information, support lifelong learning and promote reading and literacy.

4.10.1.1 P2P learning

Even though it would be safe to say that P2P computing has not yet fully matured as an educational tool (Kook, 2001: 39), indications are that the next-generation Internet will move towards the integration of some of the P2P concepts discussed in this chapter.

In the Internet Age, fast-developing technologies may have a profound impact on learning environments, particularly in the classroom (Ravenscroft, 2001). The implementation of these technological advances should encourage any-time and any-place learning that could result in learning outcomes of equal standing as those achieved in terms of the traditional-classroom format.

Statistics are available on students' learning retention as a function of the means of communication and/or the method of instruction employed by the facilitator. Following, a list of the average retention rates for the various means of communication and/or methods of instruction (Yelle, 2002):

- Lectures: an average retention rate of 5%.
- Reading: an average retention rate of 10%.
- Audio-visual modes of communication: an average retention rate of 20%.
- Demonstrations: an average retention rate of 30%.
- Discussion groups: an average retention rate of 50%.
- Practical implementation/experiments: an average retention rate of 75%.
- Teaching others, together with immediate use: an average retention rate of 80%.

These results show considerable variation between the various methods of instruction. The pattern represented by these statistics, however, also provides a useful insight, namely that P2P learning is closest to the last three results obtained, as it relies heavily on the "direct discussion", "knowledge-sharing" and "teaching others" modes of communication (Kook, 2001: 40).

4.10.1.2 Features of P2P computing contributing to P2P learning

Section 4.7, entitled "Applications utilising P2P computing", was devoted to an in-depth analysis of the application of P2P computing. This section will now be used briefly to look at its role in P2P learning:

- File-sharing: sharing of educational material without the dependence on centralised control and administration via a server.
- Instant messaging: allowing the instantaneous collaboration and communication necessary within the virtual environment of the Internet.
- Group collaboration: this feature enables both educators/facilitators and learners/students distant from each other in time and space to collaborate, thus facilitating the crucial collaborative aspect of learning.

The most obvious role that P2P computing has to play in P2P learning, however, is in a library environment. The next section will, therefore, be used to examine the possible impact of P2P computing on libraries.

4.10.2 P2P and libraries

Technology has always affected libraries, and always will. In the recent past, new technologies have, for instance, compelled libraries to improve and expand on their existing services. It is also expected that, as the 21^{st} century gathers momentum, dozens of emerging technologies will further challenge libraries at a more fundamental level, which technologies will, no doubt, include P2P networking (Guscott, 2001).

Potential ways in which P2P computing could contribute to the services rendered by libraries are as follows (Dong, 2002: 148-149):

- Personal library: an effective application would be to let the user create his/her own personal library via the Internet, finding files and sorting them in a way that would be logical and relevant, just like his/her bookshelves at home. On a P2P network, individuals would begin to obtain copies of articles, papers and even books relating to their own interests from a large number of resources. In this way, such personal libraries would have surpassed their own boundaries in helping their "owners" to locate information entirely on their own. In addition, with collaborative object lookup architecture, a personal library could use intelligent agents to fetch files, based on personal-relevance criteria. These agents would also collaborate in a P2P and cross-platform network.
- Interlibrary loans. Interlibrary loans ("ILL", for short) can be defined as "the requesting and supplying of information resources in a local and/or global fashion". Within the ILL environment, major developments with regard to the International Standards Organisation (ISO) ILL Protocol are helping to facilitate P2P computing.

The ISO ILL Protocol defines a standard for ILL systems to ILL system communication, allowing any ISO ILL-compliant ILL system to communicate with any other ISO ILL-compliant ILL system for the purposes of negotiating an ILL transaction. One of the means of communication enabled by ISO ILL is that of P2P ILL communication.

In essence, the ILL system of the Borrowing Library communicates directly with the Supplying Library for the purposes of effecting an ILL transaction. The greatest challenge for the library environment here would be to add copyright compliance to P2P computing.

- E-learning network. P2P systems are poised to revolutionise the learning process. The role of the library in this revolution would be to provide relevant and up-to-date information resources.
- Wireless services. Next-generation networks must address the issue of unique and dynamic content on occasionally connected and mobile devices. To address mobile issues in libraries also presents a promising prospect. Some libraries are offering access to the online system via wireless technology. Technically, it is already possible to build portable, anonymous, wireless, P2P file-sharing networks, by gluing together off-the-shelf technologies such as 802.11b, IP networking, HTTP services and ordinary notebook PCs (Dong, 2002: 149).

4.11 Implementation within a business

According to Adams (2000), the potential benefits to be derived from P2P networking for global corporations have tempted many technology executives into calling it "the next big thing". Companies are viewing P2P technology as an efficient and inexpensive way in which to do business, especially of the business-to-consumer and business-to-business varieties.

The business environment has evolved a number of areas in which P2P-sharing could be implemented most effectively and efficiently. Section 4.7 of this study, entitled "Applications utilising P2P computing", provided an in-depth analysis of the application of P2P computing. The next section will be used, albeit briefly, to expand on the features identified in the said section.

The following P2P computing features can be used in a business environment:

- File-sharing
- Instant messaging
- Group collaboration
- Distributed computing.

Next, each of the identified features will be discussed briefly.

- File-sharing

  The 21st century has heralded a brand-new management paradigm, as the success of a business no longer relies solely on its production, management and/or any other function *per se*, but more and more so on how effectively it obtains and implements information. To be more specific, the success of a business now relies more on the pool of knowledge shared by its employees – its human capital.

Businesses employing the more traditional client/server model merely hope that their employees would create a pool of knowledge and that they would share this in terms of the said model. Owing to human nature and a myriad of reasons that fall beyond the scope of this study, however, employees are more likely to hoard information than to share it with their colleagues. One way of counteracting this is to introduce a P2P computing network into the business, thus stimulating communication and creating a more comfortable zone for employees to share inherent knowledge of their respective fields of expertise.

- Instant messaging

  The instant-messaging feature of P2P computing will enable employees to interact in an instantaneous fashion in a bid to solve and overcome business-related problems and obstacles.

- Group collaboration

  P2P computing empowers individuals and teams to create and administer real-time and offline collaboration areas in a variety of ways, be it administered, unadministered, across the Internet or behind a firewall (P2P Working Group, 2002). In addition, P2P computing allows for online collaboration between employees who are geographically dispersed, thus decreasing the need for e-mail communication.

- Distributed computing

  Distributed computing constitutes the most innovative use of P2P technology yet, as it could aid businesses in meeting large-scale computer-processing needs (Grimes, 2001), as in the case of complex mathematical computations found in the Sciences. The ability to leverage dormant PCs for running applications requiring heavy processing reduces costs and facilitates faster completion times.

4.12 Summary

This chapter was devoted to an analysis of the following aspects: firstly, the historical foundations at the core of the P2P computing revolution and evolution. Expanding on that, the definition and aspects pertaining to P2P computing within the modern technological revolution were formulated and identified. The focus then fell on the models and related architectures associated with P2P computing. Building on the foregoing, some advantages and disadvantages were also identified. Lastly, the application of P2P and the implementation of P2P computing technology were discussed, as well as its application in business and tertiary-education institution contexts.

The next chapter will be devoted to a discussion on the results and frequencies of use of the P2P applications.