

Peer-to-peer computing

Copyright ©2006

Damir Lizdek and Dan Rundlöf.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

Chapter 1

Abstract

Peer-to-peer computing is an increasingly popular method to share resources and information among users. The most common use of peer-to-peer computing today is file sharing and there are many protocols and applications available for this purpose. The technology is also used for scientific purposes where the traditional client/server model is not capable of handling the scale of the problem.

Since most material shared in peer-to-peer file sharing networks is copyright protected the technology is legally controversial and it has been involved in several lawsuits.

Chapter 2

Preface

We are two students who study Computer Engineering at Chalmers University of Technology and we have chosen computer communication as one of our majors.

We write this report as an assignment in the course Computer Communication and Distributed Systems at Chalmers. Our examiner for the course is Philippos Tsigas at the Division of Computing Science at Chalmers.

The purpose of this report is to present peer to peer computing in an understandable and interesting way.

Contents

1	Abstract	1
2	Preface	2
3	Introduction	5
3.1	Overview of peer-to-peer computing	5
3.2	Types of peer-to-peer networks	5
3.3	Advantages of peer-to-peer networks	6
4	Different uses of peer-to-peer computing	7
4.1	File sharing	7
4.2	Bioinformatics	7
4.3	Artificial intelligence	7
4.4	Grid computing	8
5	Protocols	9
5.1	Bittorrent	9
5.1.1	How BitTorrent works	9
5.2	Fasttrack	10
5.2.1	How Fasttrack works	10
5.3	Gnutella	11
5.3.1	How Gnutella works	11
5.4	Freenet	12
5.5	DNS	12
6	Legal controversy	13
7	Anonymous peer-to-peer	15
8	Conclusion	16

9 Terminology	17
9.1 Abbreviations	17
10 References	18
A Licenses	19
A.1 GNU Free Documentation License	19

Chapter 3

Introduction

3.1 Overview of peer-to-peer computing

A peer-to-peer network is a decentralized computer network that relies mostly on the resources of equal peers instead of a few powerful servers. The peers are often connected as a ad hoc network. Such networks can be used for many things, for example sharing files, real-time data and telephony traffic.

In a pure peer-to-peer network there are no dedicated clients or servers, but instead all the peers can function as both clients and servers depending on the circumstances. A typical example for a non peer-to-peer network is an FTP server where the client and server programs both have different distinct roles where the client always initiates the transfer and the servers role is to service the clients request.

3.2 Types of peer-to-peer networks

Pure peer-to-peer:

- Peers act as both clients and server.
- There is no central server managing the network.
- There is no central router.

Hybrid peer-to-peer:

- Has a central server that keeps information on peers and responds to requests for that information.

- Peers are responsible for hosting the information (as the central server does not store files), for letting the central server know what files they want to share, and for downloading its shareable resources to peers that request it. [1]

3.3 Advantages of peer-to-peer networks

An important idea in peer-to-peer networking is that all the peers provide resources to the network. These resources can be anything from bandwidth and storage space to computing power depending on the purpose of the network. Thus if more nodes connect to the network and the demand on the system increases, the total capacity of the system also increases. This is a major advantage over traditional networks using the client/server model where additional clients could mean less resources available for each client.

Peer-to-peer networks are particularly useful when distributing large files because of the bandwidth requirement. Another advantage is that the robustness increases since there is no single point of failure in a pure peer-to-peer network. Even if it is a hybrid peer-to-peer network with a single server data can be distributed among several peers and thus making it more available. [1]

Chapter 4

Different uses of peer-to-peer computing

4.1 File sharing

File sharing is the practice of making files available for other users to download over the Internet and smaller networks. Usually filesharing implements the peer-to-peer networking model where the peers in the network are the personal computers of the users. The shared files are thus stored on the computers and are made available so that other users can download them.

Peer-to-peer file sharing is distinct from file trading in that downloading files from a peer-to-peer network does not require uploading, although some networks either provide incentives for uploading such as credits or force the sharing of files being currently downloaded.

4.2 Bioinformatics

Peer-to-peer networks have also begun to attract attention from scientists, especially those that deal with large datasets such as bioinformatics. Peer-to-peer networks can be used to run large programs that handle large amounts of data, for example tests to identify drug candidates. This is possible since peer-to-peer networks scale well. [1]

4.3 Artificial intelligence

Swarm intelligence (SI) is an artificial intelligence technique based around the study of collective behaviour in decentralised, self-organised, systems.

The expression "swarm intelligence" was introduced by Beni and Wang in 1989, in the context of cellular robotic systems. [2]

SI systems are typically made up of a population of simple agents that interact locally with each other and with their environment. Although there is normally no centralised control structure dictating how individual agents should behave, local interactions between such agents often lead to the emergence of global behaviour. Examples of systems like this can be found in nature, including ant colonies, bird flocking, animal herding, bacteria molding and fish schooling.

A peer-to-peer network can be used to model swarm intelligence where each peer in the network corresponds to a single agent.

4.4 Grid computing

Grid computing is an emerging computing model that provides the ability to perform higher throughput computing by taking advantage of many networked computers to model a virtual computer architecture that is able to distribute process execution across a parallel infrastructure.

Grids use the resources of many separate computers (peers) connected by a network (usually the Internet) to solve large-scale computation problems. Grids provide the ability to perform computations on large data sets, by breaking them down into many smaller ones, or provide the ability to perform many more computations at once than would be possible on a single computer, by dividing the labour between multiple peers. [1]

Chapter 5

Protocols

5.1 Bittorrent

BitTorrent is the name of a file sharing protocol used for peer-to-peer file distribution and it is also the name of a client application that implements the protocol, both of which were created by programmer Bram Cohen. BitTorrent is designed to widely distribute large amounts of data without incurring the corresponding consumption in costly server and bandwidth resources.

5.1.1 How BitTorrent works

BitTorrent works so that all the peers that are involved in the downloading of a file are also contributing to the distribution of the file by simultaneously making the file available for download. The sharing of a complete file is called seeding.

BitTorrent greatly reduces the load on seeders, because clients generally download the file from each other. After the initial pieces are transferred from the original seeder, the pieces are individually transferred from client to client. Because of this the original seeder only needs to send out one copy of the file to make it available for all the clients.

The BitTorrent protocol starts by breaking down the files into smaller pieces called fragments. A typical size of a fragment is a quarter of a megabyte (256 KB), though larger files can have larger fragment size. By default a 4.37GB file will have a fragment size of 4mb. Peers download missing fragments from each other and simultaneously upload those that they already have to peers that request them.

The protocol chooses the peer with the best network connections for the fragments that it is requesting. To increase the overall efficiency of the swarm, the bittorrent clients request from their peers the fragments that are most rare; in other words, the fragments that are least available, making most fragments available widely across many machines and that way avoiding bottlenecks. There is no rule that says that the fragments will be downloaded in sequence and therefore they need to be reassembled by the receiving machine. Clients will start to upload fragments to their peers before the entire file is downloaded. Sharing by a peer begins with a .torrent file, which is a small file that contains the filename and the size of the file.

Downloading with BitTorrent is easy and straightforward. Each person who wants to download the file must first start by downloading the .torrent pointer and open it in the BitTorrent application. The torrent file tells the client the address of the tracker, which, in turn, maintains a log of which users are downloading the file and where its fragments reside. For all sources available, the client will consider which blocks of the file are available and then request the rarest block that it does not yet have. This makes it more likely that peers will have blocks to exchange. As soon as the client finishes importing a block, it hashes it to make sure that the block matches what the torrent file said it should be. [3]

5.2 Fasttrack

FastTrack is another peer-to-peer file sharing protocol. It is used by several file sharing applications including Kazaa. FastTrack has been one of the most popular file sharing protocols and it was mainly used to exchange digital music files. Popular features of FastTrack are its ability to resume an interrupted download and to simultaneously download segments of one file from multiple peers increasing the overall download speed.

5.2.1 How Fasttrack works

FastTrack is a second generation peer-to-peer protocol. Second generation peer-to-peer protocols use supernodes to improve scalability. The supernode functionality is built into the client; if a powerful computer with a fast network connection runs the client software, it will automatically become a supernode, effectively acting as a temporary indexing server for other, slower clients.

In order for a user to initially connect to the network, a list of supernode IP numbers is stored in the program. The client attempts to contact these, and as soon as it finds a working supernode, it requests a list of currently active supernodes, to be used for future connection attempts. The client picks one supernode as its connection to the network and uploads a list of files it intends to share to that supernode. It also sends search requests to this supernode. The supernode communicates with other supernodes and forwards search requests to them. When a result is found it is returned to the requesting client which in turn will connect directly to the peer that has the file available and the file is transferred directly from the peer to the client using HTTP. [4]

5.3 Gnutella

Gnutella is a file sharing network that is mainly used to exchange music, films and software. It is a pure peer-to-peer network and files are transferred directly between users.

Applications that implement the Gnutella protocol connects to the Gnutella network and begin sharing files. Search queries are passed from one node to another in round-robin fashion. Gnutella clients are available for a number of different platforms.

5.3.1 How Gnutella works

To understand how the Gnutella protocol implements file sharing, imagine a large circle of nodes, who each have Gnutella client software. The first thing that takes place is that the client software finds at least one of those other nodes. Different methods have been used for this, including a pre existing list of possibly working node addresses shipped with the software, using Gwebcache sites on the web to find nodes, as well as using IRC to find nodes. Chances are that at least one other node (call it B) will be found. Once it has connected, node B will send node A its own list of working nodes. Node A will try to connect to the nodes it was shipped with, as well as nodes it receives from other nodes, until it reaches a certain quota, usually user-specifiable. It will only connect to that many nodes, but it keeps the nodes it has not yet tried. (It discards ones that it tries but did not work.)

Now, when user A wants to do a search, it sends the request to each node it is actively connected to. It is possible that some of them will no longer work, in which case user A tries to connect to the nodes it has saved as backups. The

number of actively connected nodes for user A is usually quite small (around 5), so each node then forwards the request to all the nodes it is connected to, and they in turn forward the request, and so on. In theory, the request will eventually find its way to every user on the Gnutella network. As the size of the Gnutella network has grown and its developers have fought excess traffic consumption, global searchability of the network has diminished.

If a search request turns up a result, the node that had the result contacts the searcher either directly or indirectly. If the node that sent the search request is not protected by a firewall, the node with the result directly returns the result. If the node that sent the search request is protected by a firewall, then the result is (indirectly) routed back along the route the search was received on. After the result is returned, they negotiate the file transfer and the transfer proceeds. If more than one copy of the same file is found, the searcher can perform a "swarm" download - download pieces of the file from different nodes. This results in increased download rates.

Finally, when user A disconnects, the client software saves the list of nodes that it was actively connected to, and that it was keeping as a backup, for use next time it connects. [5]

5.4 Freenet

Freenet is a decentralized network that uses a peer-to-peer protocol that provides anonymity for its users. The idea behind it is to provide electronic freedom of speech. Freenet uses the bandwidth and resources of the peers in the network to allow them to anonymously publish or retrieve information amongst them.

Currently, Freenet is unable to handle dynamic content, such as content that utilizes databases and scripting. According to the Freenet Project group, such tradeoffs are expected since Freenet's primary goals are neither ease-of-use nor performance. Unlike other peer-to-peer networks, Freenet is primarily intended to combat censorship and allow people to communicate freely and with near-total anonymity.

Freenet can be thought of as a large storage device. When you store a file in it, you receive a key which can be used to retrieve the file. When you supply Freenet with a key, it will return the appropriate file (if it is located). The Freenet storage space is distributed among all connected nodes on Freenet. [6]

5.5 DNS

While peer-to-peer computing is mostly used for file sharing it can also be used for other purposes such as the Domain Name System protocol. The goal of the protocol is to provide a name translation service from domain names to IP addresses. DNS is not a pure peer-to-peer protocol since it relies on some non-peer elements.

Chapter 6

Legal controversy

The most common use for peer-to-peer computing today is filesharing. Since the majority of the material that is downloaded is protected by different copyright laws there is an obvious conflict of interest between individuals sharing and downloading files and the copyright holders.

This has led to the fact that peer-to-peer computing has been targeted in lawsuits and the technology is still controversial. In the Sony Corp. v. Universal City Studios case “the Betamax decision” states that copying technologies are not inherently illegal, if there are legal uses of the technology. This decision applies to peer-to-peer networks since they can be used to distribute that are not copyrighted as well. [6]

As we pointed out earlier the majority of the files shared in peer-to-peer networks are various music and video files that are copyrighted. Sharing of these files is illegal in most cases and that is the main reason for the legal controversy that surrounds peer-to-peer computing. This is also the reason why global organisations such as RIAA and the MPAA now spend large amounts of money attempting to lobby governments to legislate to protect copyright holders and their work.

In spite of the Betamax decision, these global organizations often target the peer-to-peer networks themselves and try to illegalize them. An example of this is when RIAA filed a lawsuit against Napster, the most popular music sharing application at the time. This led to the shutdown of the Napster network and bankruptcy of the company.

As the actions to protect the copyrights of media companies expand, the file sharing technology has quickly adapted and is constantly developing making it much more difficult to dismantle and shutdown. This instead has caused

the users to become targets of the lawsuits since the technology in itself might be legal but certain uses of it might be illegal.

In the European Union, the 2001 EU Copyright directive, which implemented the 1996 WIPO treaty prohibits peer-to-peer, claiming it is a violation of the directive. However, not all EU member states have implemented the directive in their national legislation. [8]

Chapter 7

Anonymous peer-to-peer

An anonymous peer-to-peer computer network is a special kind of peer-to-peer network where the users of the network are anonymous by default. The primary difference between regular and anonymous networks is in the routing methods of the respective networks.

The peer-to-peer file sharing community's interest in anonymous networks has increased rapidly in recent years for many reasons. One of them is to avoid censorship and another is to circumvent the legal issues that are associated with file sharing.

Many claim that the anonymous networks will lead to an increase of illegal file sharing and some even suggest that they facilitate terrorism. Others counter that the potential for illegal uses should not prevent the technology from being used for legal purposes and that the presumption of innocence must apply.

Chapter 8

Conclusion

Peer-to-peer computing is a powerful technology that is continuously evolving and finding new uses. We believe that the technology will be increasingly important as new areas adopt it.

File sharing is the most popular area of peer-to-peer computing today. We think that this will be the case for many years to come and it is therefore important to solve the legal issues surrounding file sharing in particular and peer-to-peer computing in general. We do not believe that illegalizing the use of file sharing applications and filing law suits against the users is the solution to the problem. Instead we would like to see a way to legalize file sharing by some sort of payment arrangement that makes sure the creators of copyrighted material are paid for their work.

Chapter 9

Terminology

9.1 Abbreviations

- **DNS** Domain Name System.
- **FTP** File Transfer Protocol.
- **GFDL** GNU Free Document License.
- **HTTP** Hyper Text Transfer Protocol.
- **IP** Internet Protocol.
- **MPAA** Motion Picture Association of America.
- **P2P** Peer-to-peer.
- **RIAA** Recording Industry Association of America.
- **WIPO** World Intellectual Property Organization.

Chapter 10

References

[1] <http://en.wikipedia.org/wiki/P2p>

[2] G. Beni and J. Wang, Swarm Intelligence in Cellular Robotic Systems, Proceedings of the NATO Advanced Workshop on Robots and Biological Systems, Il Ciocco, Tuscany, Italy, 1989.

[3] <http://www.bittorrent.org/protocol.html>

[4] <http://cvs.berlios.de/cgi-bin/viewcvs.cgi/gift-fasttrack/giFT-FastTrack/PROTOCOL?rev=HEAD&content-type=text/vnd.viewcvs-markup>

[5] http://rfc-gnutella.sourceforge.net/src/rfc-0_6-draft.html

[6] <http://freenet.sourceforge.net/index.php?page=whatis>

[7] http://en.wikipedia.org/wiki/Sony_Corp._of_America_v._Universal_City_Studios%2C_Inc

[8] <http://art.ntu.ac.uk/liveart/issues/Chapter11.htm>

Appendix A

Licenses

A.1 GNU Free Documentation License

Version 1.2, November 2002

Copyright ©2000,2001,2002 Free Software Foundation, Inc.

51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The **”Document”**, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as **”you”**. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A **”Modified Version”** of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A **”Secondary Section”** is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The **”Invariant Sections”** are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The **”Cover Texts”** are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A **”Transparent”** copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image

format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "**Opaque**".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "**Title Page**" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "**Entitled XYZ**" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "**Acknowledgements**", "**Dedications**", "**Endorsements**", or "**History**".) To "**Preserve the Title**" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.

- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or

any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright ©2006 Damir Lizdek & Dan Rundlöf. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this:

with the Invariant Sections being **Abstract, Preface, Introduction, Different uses of peer to peer computing, Protocols, Legal controversy, Anonymous peer to peer computing, Conclusion, Terminology, References**, with the Front-Cover Texts being **Peer to peer Computing**.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.