

University of Southern Queensland
Faculty of Engineering & Surveying

**Design and Implementation of a Network Address
Translator**

A dissertation submitted by

K-J. Beasley

in fulfilment of the requirements of

ENG4112 Research Project

towards the degree of

**Bachelor of Engineering (Computer Systems Engineering) / Bachelor
of Information Technology (Applied Computer Science)**

Submitted: October, 2004

Abstract

A continuously increasing demand for Internet Protocol (IP) Addresses was something that was not considered at the time when the Internet was first designed. The argument was actually quite the opposite and most experts pooh-pooed the idea of the internet ever growing to beyond 100,000 networks. However, the 100,000th network was connected to the internet in 1996 (Tanenbaum 2003). The Explosive growth of the Internet has resulted in a shortage of the number of available IP Addresses. As this growth continues the shortage will increase and a new form of Internet Addressing will need to be established. The current form, IPv6 has been under development for some time now and has not gathered wide industry support. Obviously a temporary solution must be established to overcome the shortage of IP Addresses in the immediate future until permanent solutions can be achieved.

On the other hand the TCP Protocol was established as an end-to-end connection for reliable communication and makes use of its own 16-bit port number. This allows for up to 65,535 unique port numbers for TCP communication. Most hosts never maintain 65,535 end-to-end connections and this allows for a technology called Network Address Translation (NAT) to save on the number of IP Addresses required on the Internet by multiplexing many IP Sources onto one or more IP Addresses using unique TCP port numbers for each data stream.

The ultimate aim of this project is to produce a small prototype Network Address Translator and discuss further improvements necessary for its use in a production environment.

University of Southern Queensland
Faculty of Engineering and Surveying

ENG4111/2 <i>Research Project</i>
--

Limitations of Use

The Council of the University of Southern Queensland, its Faculty of Engineering and Surveying, and the staff of the University of Southern Queensland, do not accept any responsibility for the truth, accuracy or completeness of material contained within or associated with this dissertation.

Persons using all or any part of this material do so at their own risk, and not at the risk of the Council of the University of Southern Queensland, its Faculty of Engineering and Surveying or the staff of the University of Southern Queensland.

This dissertation reports an educational exercise and has no purpose or validity beyond this exercise. The sole purpose of the course pair entitled “Research Project” is to contribute to the overall education within the student’s chosen degree program. This document, the associated hardware, software, drawings, and other material set out in the associated appendices should not be used for any other purpose: if they are so used, it is entirely at the risk of the user.

Prof G Baker

Dean

Faculty of Engineering and Surveying

Certification of Dissertation

I certify that the ideas, designs and experimental work, results, analyses and conclusions set out in this dissertation are entirely my own effort, except where otherwise indicated and acknowledged.

I further certify that the work is original and has not been previously submitted for assessment in any other course or institution, except where specifically stated.

K-J. BEASLEY

00110217980

Signature

Date

Acknowledgments

I would like to thank all those who have supported my endeavors over the past five years. Special thanks goes to all the staff from the Faculty of Engineering and Surveying and the Department of Maths and Computing at the Faculty of Science. Students are incredibly grateful for the time you take to pass on your knowledge and experience.

I would like to offer special thanks to my supervisor, Dr John Leis. After debugging a single line of code for what seemed like a month, losing all motivation and wondering why I ever suggested this project topic, it was often his support and desire to see the project succeed that positioned me back on track.

Sincere appreciation to my family. I have no idea where I would be right now if it were not for my Mum constantly looking over my shoulder checking if my project was up to date and my assignments completed. It was Dad's seemingly infinite technical knowledge that convinced me to study Engineering. Sadly this knowledge is much diminished due to a recent debilitating stroke. Lastly to my darling sister, she may be public enemy number one for telling Mum that I worked on project/assignments when indeed I played computer games, however her love and support during this time is greatly appreciated. The completion of this project and my successful graduation is testimony to your contributions and continued support.

Thankyou All.

K-J. BEASLEY

University of Southern Queensland

October 2004

Contents

Abstract	i
Acknowledgments	iv
List of Figures	xi
List of Tables	xiii
Chapter 1 Introduction	1
1.1 Overview of the Dissertation	3
Chapter 2 Network Reference Models	4
2.1 Chapter Overview	4
2.2 Networking History	4
2.3 The Open Systems Interconnection Reference Model	6
2.3.1 The Physical Layer	7
2.3.2 The Data Link Layer	7
2.3.3 The Network Layer	7

CONTENTS	vi
2.3.4 The Transport Layer	8
2.3.5 The Session Layer	9
2.3.6 The Presentation Layer	9
2.3.7 The Application Layer	9
2.4 The TCP/IP Reference Model	9
2.4.1 The Host-to-Network Layer	10
2.4.2 The Internet Layer	11
2.4.3 The Transport Layer	11
2.4.4 The Application Layer	12
2.5 Chapter Summary	13
 Chapter 3 Design Specification	 14
3.1 Chapter Overview	14
3.2 Design Methodology	14
3.2.1 Extreme Programming	14
3.2.2 Scrum	15
3.2.3 Feature Driven Development	15
3.2.4 STEPWISE	16
3.2.5 Rational Unified Process	16
3.2.6 Waterfall Models	17
3.3 Programming Language	18

CONTENTS	vii
3.3.1 Visual Basic.NET	19
3.3.2 ASP.NET	19
3.3.3 C	19
3.3.4 C++/Visual C++.NET	20
3.3.5 Java	21
3.3.6 C#.NET	21
3.4 Chapter Summary	23
Chapter 4 Internet Protocol	24
4.1 Chapter Overview	24
4.2 IP Functionality	24
4.2.1 Looking at an IP Header	25
4.2.2 Fragmentation	29
4.2.3 IP Addresses	30
4.2.4 Subnets	31
4.2.5 Classless InterDomain Routing	31
4.3 Network Address Translation	32
4.3.1 Overview of NAPT	33
4.3.2 Address Binding	34
4.3.3 Address Unbinding	34
4.3.4 Header Manipulation	34

CONTENTS	viii
4.3.5 Incremental Checksum Adjustment	35
4.3.6 ICMP error packet modifications	35
4.3.7 FTP Support	35
4.3.8 Using IP Options	36
4.3.9 Recommendations for Private IP Address Range	36
4.3.10 Privacy and Security	36
4.3.11 Fragmented Packets	37
4.4 IPv6	37
4.5 Chapter Summary	38
Chapter 5 Transmission Control Protocol	39
5.1 Chapter Overview	39
5.2 TCP Functionality	39
5.2.1 Addressing	42
5.2.2 Reliability	42
5.2.3 Congestion Control	43
5.2.4 Connection Management	44
5.3 Data Connections	44
5.3.1 Three Way Handshaking	44
5.3.2 Simultaneous Open	45
5.3.3 Active Close	45

CONTENTS	ix
5.3.4 Passive Close	46
5.4 User Datagram Protocol / Real-Time Transport Protocol	47
5.5 Chapter Summary	48
Chapter 6 Existing Network Address Translators	49
6.1 Chapter Overview	49
6.2 Windows	50
6.2.1 NAT32E	50
6.2.2 BrowseGate 3 NAT/Proxy server and firewall	50
6.3 Linux	51
6.3.1 IP Masquerading	51
6.3.2 IP Tables	51
6.4 Chapter Summary	53
Chapter 7 Network Address Translator Implementation	55
7.1 Chapter Overview	55
7.2 C#.NET Basics	55
7.3 Using Sockets	59
7.3.1 Application Programming Interface	59
7.3.2 Windows Sockets	59
7.3.3 Advanced Socket Control	60

CONTENTS	x
7.4 Putting it all Together	61
7.4.1 Pseudocode	62
7.4.2 A Windows Service	62
7.5 Chapter Summary	65
Chapter 8 Conclusions and Further Work	66
8.1 Achievement of Project Objectives	66
8.2 Further Work	68
References	69
Appendix A Project Specification	71
Appendix B Project Source Code	73
B.1 NATService.cs	74
B.2 RawSocket.cs	81
B.3 BidirHashtable.cs	105
B.4 ProjectInstaller.cs	109
B.5 NATControl.cs	113

List of Figures

2.1	(a) Structure of a Switching office. (b) Structure of the telephone system. (c) Baran's proposal for a distributed switching system. (adapted from (Baran 1964)).	5
2.2	The Open Systems Interconnection Reference Model. (adapted from (Day & Zimmermann 1983)).	6
2.3	The TCP/IP reference model.	10
4.1	The IPv4 (Internet Protocol) header. (adapted from (Tanenbaum 2003)). . . .	25
4.2	An Example of connecting two networks with differing MTU values. (adapted from (Feit 1998)).	30
5.1	The IPv4 (Internet Protocol) header. (adapted from (Tanenbaum 2003)). . . .	40
5.2	Combining TCP and IP to encapsulate data.	40
5.3	TCP connection establishment.	45
5.4	TCP Simultaneous Open. (adapted from (Tanenbaum 2003))	46
7.1	A simple C#.NET form.	56
7.2	The result of Hello World Code execution on the C#.NET form.	58

7.3 The Visual Studio Development Environment.	60
--	----

List of Tables

4.1	Contents of a real IP packet.	28
-----	---------------------------------------	----

Chapter 1

Introduction

Explosive growth of computer networks, in particular the Internet has seen the Internet become an integral part of everyday life. Many of the tasks traditionally left for the Mail network are now being done via e-mail. Phone conversations are increasingly being replaced by Internet Chat and slowly voice chat and webcams are entering the market while available bandwidth is making this technology viable.

There is a significant trend towards internet connectivity for devices which previously would never have been considered relevant to the Internet. For example Internet Refrigerators and Air Conditioners which can be activated from a remote device are becoming popular. As development continues more and more devices will use an Internet connection as part of their operation.

The Internet currently operates based on two important protocols collectively referred to as TCP/IP. These two protocols are actually the Transmission Control Protocol and Internet Protocol. Unfortunately the designers of IP overlooked the commercial viability of the Internet and suspected it would never become anything more than a research network connecting universities and a few other large companies such as Military Research and Development.

This line of thought lead to the development of the Internet Protocol utilizing an Addressing system consisting of 32 bit addresses. That is addresses consist of 32 binary

numbers each of which can only be on (1) or off (0). In an ideal world this offers $2^{32} = 4,294,967,296$ possible IP Addresses. Given that there are in excess of 6 billion people in the world this number is never going to survive in the long term when a very large proportion of the population will hopefully be online. In addition many people use more than one IP Address, for example if they have more than one computer connected directly to the internet or have a computer at work and at home which may both be connected at the same time. Finally we do not live in the ideal world and IP addresses are wasted both through requirements for divisions and wastage due to over-allocation or private requirements.

The obvious solution to this would be to allow the IP Address to be larger, maybe double or quadruple its current size. Such a protocol is being developed and has already been implemented in some areas. Unfortunately there has not been a great deal of industry support for the new protocol and some believe it may lose all levels of support and be forgotten before it is implemented. However if this happens we will still be left to face the problem of how to spread some four billion IP Addresses across the globe in a fair and equal manner. In addition some companies and universities who have purchased large ranges of the IP address space will not likely give up their range unless presented with sufficient financial incentive to do so. As the remainder of a particular resource decreases its value normally increases and this certainly could be the case with IP Addresses.

Unfortunately the lack of IP Addresses is a real difficulty affecting people across the Internet at this time. The problem is not going to wait for a solution to be developed, implemented and tested. Therefore alternatives must be developed quickly, must require little testing and must be reliable. One such solution is called Network Address Translation (NAT) and is already implemented in many forms. In Microsoft Windows 98 and above the solution is commonly known as Internet Connection Sharing (ICS). However ICS is a poor solution to the problem which works in some cases but excludes any sort of special protocols such as Video Conferencing (NetMeeting) and active File Transfer Protocol (FTP). Due to the ICS code forming an integral part of the Operating System it cannot be reverse engineered or modified and has very few security features which has allowed other products to enter the market.

The aim of this project is to develop the basics of a new type of Network Address Translator. Ideally the NAT will undergo further development after the initial project is complete and will contain security features and policies rigid enough to satisfy even the most security conscious Network Administrator while having the flexibility to be used by even the most application intensive home Internet user regardless of which applications they may wish to use.

1.1 Overview of the Dissertation

This dissertation is organized as follows:

Chapter 2 reviews the beginning of the Internet and some of the network models used in designing the communications protocols used today.

Chapter 3 discusses the development of a Network Address Translator (NAT), the programming methodology followed and the Programming Languages available for Implementation.

Chapter 4 examines the Internet Protocol (IP) used in network communications. The Internet Protocol forms part of the network hierarchy and is the first layer involved in Network Address Translation.

Chapter 5 details the Transmission Control Protocol (TCP) which completes the TCP/IP Protocol Suite used on the Internet. The use of Transmission Control Protocol features in Network Address Translation are also discussed.

Chapter 6 critically examines existing Network Address Translators including popular features for Home and Business users and cost to purchase. New Network Address Translator features which may be well received by Home and Business users are also discussed.

Chapter 7 introduces the C#.NET programming language and development environment. The most important part of this project, communication sockets, are introduced including advanced features required by this project.

Chapter 8 concludes the dissertation and suggests further work in the area of 'z'.

Chapter 2

Network Reference Models

2.1 Chapter Overview

Looking at the history of Network Design, many networks were mainly hardware oriented with the software as an afterthought. This strategy is no longer suitable for today's high speed networking interfaces. This chapter will examine the software structuring in some detail.

2.2 Networking History

At the height of the Cold War in the late 1950's one line of thought was the vulnerability of the telephone network to Nuclear War. (Baran 1964) Referring to Figure 2.1(b) reveals that the destruction of a few key points could fragment the telephone network into small isolated islands.

Around 1960, the Department of Defense (DoD) awarded a contract to RAND Corporation for the development of a solution to this vulnerability. Paul Baran, an employee of RAND Corporation, developed a proposed solution depicted in Figure 2.1(c). Unfortunately When the DoD took the idea to the U.S. national telephone provider **AT&T** the idea was dismissed as a concept which could not be constructed. It is believed that **AT&T**

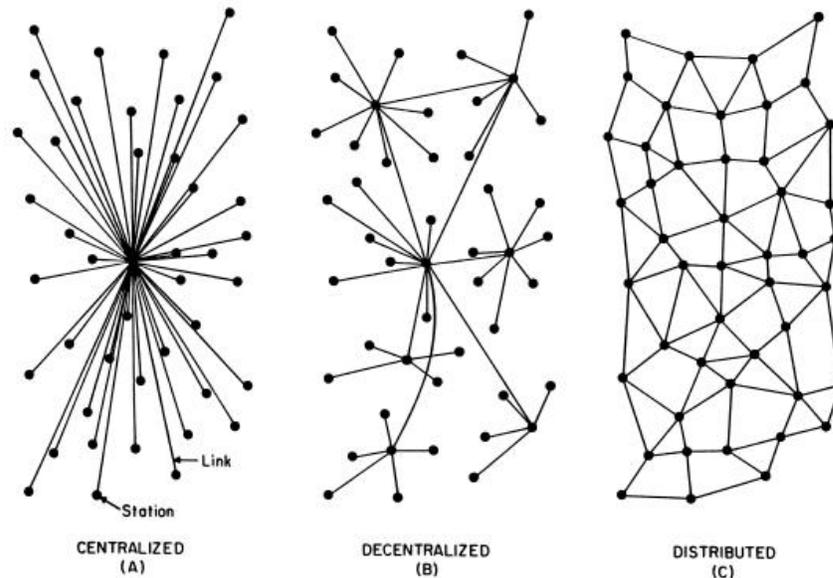


Figure 2.1: (a) Structure of a Switching office. (b) Structure of the telephone system. (c) Baran's proposal for a distributed switching system. (adapted from (Baran 1964)).

did not actually want to admit that Paul Baran had succeeded in developing a network concept where AT&T had failed, effectively dooming the idea. (Tanenbaum 2003)

In 1967 the Advanced Research Projects Agency director, Larry Roberts, turned the sights back onto networking. He worked with Wesley Clark who again suggested a packet-switched subnet communicating via routers. Roberts presented a vague paper on the packet-switching idea at the Symposium on Operating System Principles (SIGOPS) in Gatlinburg. (Roberts 1967) A similar paper at the conference described a system that had not only been designed, but actually implemented at the National Physics Laboratory in England. In 1968 BBN, a consulting firm in Cambridge, Massachusetts was awarded a contract to build what became known as the ARPANET (Tanenbaum 2003)

In the early 1980's ARPANET protocols were eventually replaced by the Transmission Control Protocol/Internet Protocol (TCP/IP) which will be discussed separately in 4 and 5. Several contracts were also awarded to BBN and the University of California at Berkeley forming the Berkeley UNIX company. Berkeley students wrote a program interface for networking called Berkeley Sockets (or simply sockets) and developed

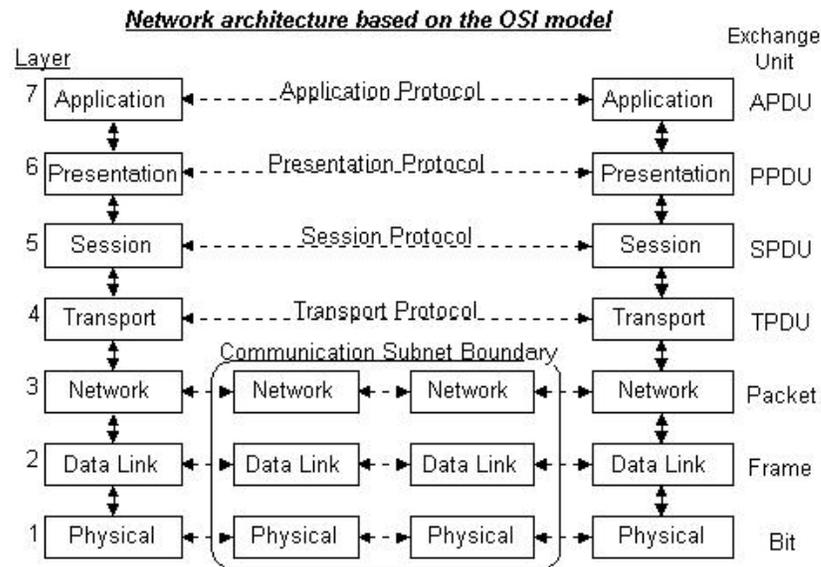


Figure 2.2: The Open Systems Interconnection Reference Model. (adapted from (Day & Zimmermann 1983)).

many applications, utilities and management programs to ease the burden of Network Administration. (Tanenbaum 2003)

2.3 The Open Systems Interconnection Reference Model

The Open Systems Interconnection (OSI) model is shown in Figure 2.2. It was based on a International Standards Organization (ISO) proposal aimed at international standardization of networking protocols.

The OSI Model was defined using five basic principles as follows:

1. A layer should add a level of abstraction to the communications architecture.
2. Each layer should provide well defined functionality.
3. The functionality of each layer should work towards an international standard.
4. Minimal information should flow across the boundaries between layers, particularly control information.

5. There should be sufficient layers that the distinct functionality of each layer is not compromised however few enough layers that the architecture is not unwieldy.

2.3.1 The Physical Layer

This layer is concerned with transmitting raw bits over a communication channel. It is concerned with issues such as timing, the standards of the physical interface, what constitutes the correct receipt of a on or off bit and generally anything to do with the electrical or mechanical characteristics of the medium used for transmission.

In recent time the physical layer has undergone significant changes due to new technology in use for physically providing the data connection. Wireless Networking is an example of one entirely new physical mediums now used in computer networks. (Day & Zimmermann 1983)

2.3.2 The Data Link Layer

The data link layer is concerned with synchronization, reliability and framing which refers to organizing each chunk of data into a packet. Frames are ordered to prevent data from arriving in the wrong order and may be acknowledged in the case of a reliable service. However this service is not concerned with ensuring reliable data streams and frames may still be corrupted or not delivered.

In the data link layer on a broadcast based network an additional issue is addressed; how to control access to the shared channel. A special sub-layer which does not form part of the OSI Model deals with this issue and is called the medium access control sub-layer. (Tanenbaum 2003)

2.3.3 The Network Layer

The main task of the network layer is to control operation of a subnet. The key design issue is routing packets from source to destination. Routing can be based on various

methods, ranging from static tables (which form a core part of the router and only change if a major reconfiguration is detected) through dynamic session based (changing when a connection is established or closed) to highly dynamic (determining different routing information for each packet based on current network load).

The Network Layer is also responsible for congestion control on the local subnet, to provide any Quality of Service controls (delay, transmit time, jitter, etc) required on the subnet and to allow the interconnection of heterogeneous networks including packet fragmentation.

Broadcast based networks often use a very small Network Layer or may not contain this layer at all. (Tanenbaum 2003)

2.3.4 The Transport Layer

The transport layer accepts data from the upper layer services, splits the data into smaller units, if necessary, and passes the chunks to the network layer. It is responsible for error detection and overall sequencing to ensure the ordering of messages is not changed in the transmission. The most important aspect of the Transport Layer is to shield the upper layer services from inevitable changes to the networking hardware.

The transport layer offers multiple types of services to the session layer. The most common service is an error-free point-to-point protocol. However other options such as best-effort transmission or real time (approximately) transmission may also be available. Obviously if the underlying layers offer broadcast and multicast services these will likely also be offered to the session layer.

The transport layer is the first end-to-end layer. Lower layers operate between neighboring machines or routers that form part of the network. The transport layer only operates on the source and destination machines which carry on the conversation. (Day & Zimmermann 1983)

2.3.5 The Session Layer

The basic function of the network layer is to allow users on different machines to establish sessions between them. The sessions offer special functionality such as dialog control (taking turns to transmit and receive), token management (preventing two machines from accessing the same area of memory or from updating the same information at the same time) and download resume features (Internet Explorer resuming half-way through a download even after a disconnection or crash). (Tanenbaum 2003)

2.3.6 The Presentation Layer

The presentation layer defines the syntax and semantics of information exchange. For example transmitting plain text in American Standard Code for Information Interchange (ASCII) or Unicode format. This includes dealing with differing methods of storing information and agreeing on a standard during transmission. A simple example might be the different methods of storing a date between America (12/31/2004) or Australia (31/12/2004). (Tanenbaum 2003)

2.3.7 The Application Layer

The application layer contains the range of higher-level protocols used by users of the Internet. These protocols include features such as file transfer, electronic mail, news servers and chat services. One of the most common protocols used for delivery of almost all internet web pages is HyperText Transfer Protocol (HTTP) which defines how to request and receive pages written in HyperText Markup Language (HTML). (Day & Zimmermann 1983)

2.4 The TCP/IP Reference Model

The TCP/IP Reference Model was designed in response to the need for the seamless interconnection of multiple networks. A common misconception is that the TCP/IP

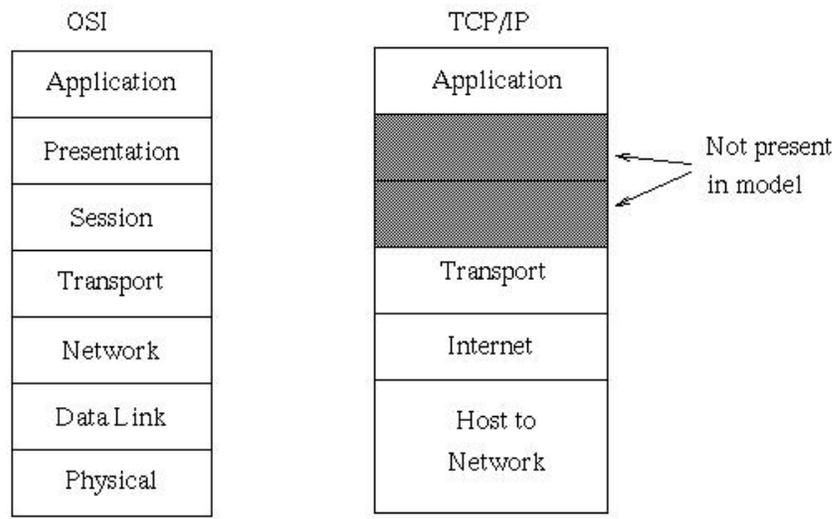


Figure 2.3: The TCP/IP reference model.

Reference Model was designed to smooth over some issues found in the OSI Reference Model. This is not true as the TCP/IP Reference Model was defined several years before the introduction of the OSI Reference Model (Cerf & Kahn 1974). The TCP/IP Reference Model, as shown in Figure 2.3, was designed mostly to satisfy Department of Defense requirements that end-to-end connections remained intact as long as the source and destination machines were functioning. This relied on the assumption that a functioning route between the source and destination existed, however, the main idea was that the actual path this connection followed could change in response to individual transmission links being decommissioned. As a result the TCP/IP Reference Model did not define anything below the Internet layer in detail. (Tanenbaum 2003)

2.4.1 The Host-to-Network Layer

The host-to-network layer is the great void left below the Internet layer. Most references on the TCP/IP reference model do not discuss this layer, however, it is included here for completeness. The main concept here is that each host must connect to the network and this will involve some protocol to encapsulate the Internet Layer, however the details of this protocol are not covered.

2.4.2 The Internet Layer

The internet layer is a connectionless internetwork that forms the basis of the TCP/IP Reference Model. Its function is to allow packets from any network to travel independently to a destination which may be separated from the packet source by many different networks. The ordering of packets may be altered during transit and each packet may follow a completely different path to the destination. This concept of following a pathway to the correct destination is known as routing.

The internet layer defines the protocol known as the Internet Protocol (IP) and associated format for an IP Header. Obviously Routing and congestion avoidance are the major issues at this level which leads to the association between the TCP/IP internet layer and the OSI network layer. (Cerf & Kahn 1974)

2.4.3 The Transport Layer

The transport layer is designed to allow a source and destination entities to undertake a conversation. The functionality here is virtually the same as in the OSI Model, the difference being that the TCP/IP Reference Model actually defines two end-to-end transport protocols. These protocols are called Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).

TCP is a reliable, connection oriented protocol which means it requires an end-to-end connection to be established and maintained. This is achieved by requiring an acknowledgement for each data segment or packet. The source is limited in how many packets can be transmitted before waiting for acknowledgements to be received. The main concept is to allow a byte-stream to be delivered from source to destination without error or corruption. TCP splits the byte-stream into fragments or discrete messages, adds error checking, sequencing and flow control information and passes the packet to the internet layer.

UDP is a unreliable, connectionless protocol meaning it does not require an end-to-end connection and could continue sending a flood of packets to a destination even

though the path may have become unavailable. Packets are not acknowledged and may arrive corrupted or may not arrive at all. Packets may also not arrive in the same order as they are sent. The purpose of providing this service is to allow application developers to implement their own sequencing and checksum's or if the application warrants, to exclude such features completely. Some applications which generally do not need sequencing or reliability through checksums are real-time voice and video transmission. (Tanenbaum 2003)

2.4.4 The Application Layer

Another void is found between the transport and application layers. The OSI Models session and presentation layers were not perceived as necessary by the developers of the TCP/IP model and were not included. Fortunately experience with the OSI Model has shown that the session and presentation layers are practically of no use to many applications.

The application layer contains all the higher-level protocols. The early internet protocols included virtual terminal (TELNET), File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP) and HyperText Transfer Protocol (HTTP). A number of additional protocols have been developed for real-time chat and voice over IP. Continuing protocol development is expected as system integrators find new and interesting ways of using the internet to make life easier. (Cerf & Kahn 1974)

2.5 Chapter Summary

The internet came from fairly humble beginnings as a small connection of four computers to form the first version of the ARPANET but not before major troubles were overcome. Developing a packet-switched subnet was something that had not been attempted and most telecommunications providers did not like the idea of some young hot shot researcher telling them how to construct their networks. However once the ARPANET began to grow, adopted more scalable protocols and became used by a large range of people the concept of a internet became possible.

A large effort began to standardize how the Internet would work. Two models were developed to address this issue. The OSI model is more generic and can be applied to almost any form of network. Years of experience have shown that the number of layers in the OSI model are slightly excessive and some layers are generally not used in Internet applications. The TCP/IP model used less layers but is more specific to the internet and can be fairly abstract when referring to lower level network features and interface. As the concepts behind software development focus more on semantics and syntaxes of the language or protocol the OSI model may gain popularity again, however, at the moment the TCP/IP model is generally most relevant to internet development.

The constant development of new protocols is a constant challenge to developers focusing on TCP/IP as new protocols may not always let TCP and IP handel the issues of routing, addressing and reliability themselves. An Early example of this was FTP where the IP address and TCP Port was embedded in the FTP data stream. As application developers develop new ways of using the Internet it is important to observe the layering of the Internet and try to avoid breaching these layers when designing new software and standards. This can particularly be a problem for applications such as Network Address Translators where failure to correct a reference to the IP address or TCP port in packet header or data can corrupt the entire process of translation.

Chapter 3

Design Specification

3.1 Chapter Overview

Many software development projects have been known to incur extensive and costly design errors. The most expansive errors are often introduced early in the development process. This underscores the need for better requirement definition and software design methodology. Software design is an important activity as it determines how the whole software development task would proceed including the system maintenance. The design of software is essentially a skill, however, it usually requires a structure which will provide a guide or a methodology for this task.

3.2 Design Methodology

3.2.1 Extreme Programming

Extreme Programming is a deliberate and disciplined approach to software Development. It has been developed over a period of about 8 years and has proven successful in companies of various sizes.

Extreme Programming is oriented towards customer satisfaction. It aims to deliver the

required software on time even when changing requirements complicate the process. Team work is central to the methodology in order to achieve this goal. (Wells 2003)

Despite these advantages, Extreme Programming is not a good choice in this project because the customer who requires the product is also the programmer, therefore changing requirements are not likely to occur. There is also no development team involved in this project so the use of GroupWise development would be a waste of effort.

3.2.2 Scrum

Scrum is an agile, lightweight process used in Product Development, particularly control and management of software projects. Scrum focuses on traditional iterative, incremental programming methods while wrapping existing engineering methodologies such as Extreme Programming and Rational Unified Process to allow agile development and simple implementation.

Scrum significantly decreases development time and has faster benefit implementation while allowing adaptive, empirical system development. (Advanced Development Methods Inc 2004)

Unfortunately Scrum is a highly commercialised development process which requires employment of a certified ScrumMaster or participation in a two day course to become a ScrumMaster. Most of these courses are only available in the United States and are financially expensive. Scrum will not be implemented as part of this project.

3.2.3 Feature Driven Development

Feature Driven Development (FDD) is a process of software development aimed at delivering requested or required features in the shortest possible time period. After the overall project is identified and a feature list is created, each identified feature is fully designed and then implemented into the system.

FDD allows a software development team to remain highly focused and greatly in-

creases production and improves team spirit by delivering entire fully featured prototypes throughout the development cycle.

Any software development suffers from exponential development times. As the project nears completion, the amount of work completed decreases for the same amount of time spent on development. FDD suffers greatly from this problem because each feature suffers from exponential development times.

FDD is particularly common in projects which are in trouble and have deadlines and milestones to be met. By focusing only on critical required features the project can often be saved. FDD will only be used in this project if development falls behind schedule.

3.2.4 STEPWISE

STEPWISE is a software development process designed to overcome limitations of the ISO 10303-11 EXPRESS model by automating software development. EXPRESS is used to represent product and process data in standard data stores, to increase data value and decrease data management costs.

STEPWISE features an enhanced architecture to support automation of EXPRESS for implementing high-level procedural interfaces, storage representations and interchange formats. (Kahn 2000)

STEPWISE is designed for high-level applications and is probably more suited to 4th and 5th level languages such as Structured Query Language. Despite its improvements over C++, C#.NET is still a 3rd level language and is not particularly suited for STEPWISE.

3.2.5 Rational Unified Process

The Rational Unified Process (RUP) was developed by the same people who originally created Unified Modeling Language (UML). UML is a single complete notation for

describing object models and is extensively used in Software Engineering. RUP is a software development process providing a framework that can be used to describe specific development processes.

The essence of RUP is iteration and RUP was developed with the goal that each iteration ends in a deliverable (prototype, fully featured class, etc). RUP involves extensive Risk Management, particularly of the risk that development will fall behind schedule. RUP acknowledges that project plans do not define what will be produced, but a statement of how to manage risk. A plan of action will inevitably fail while a plan of contingencies will eventually succeed. (Sharon 1999)

However as a major part of the Engineering Program, this project cannot be allowed to fail. Some aspects of the program development may be optional, provided that the overall deliverables are provided. For this reason the contingencies are somewhat limited and not particularly suited to RUP.

3.2.6 Waterfall Models

The waterfall model was originally developed as a series of discontinuous phases involving Conception, Requirements, Architectural Design, Detailed Design, Coding and Development and Testing and Implementation.

Several variations on this system interpose their own advantages and weaknesses into the model. These variations include the Spiral Model, Modified Waterfall Model, Evolutionary Prototyping, Code-and-Fix, Staged Delivery and Evolutionary Delivery.

- The Spiral Model breaks a software project up into mini-projects, each addressing a major risk. This ensures that total project risk is inversely proportional to cost at each step in the development process.
- The Modified Waterfall Model is potentially the same as the Waterfall Model, however it is not done in Discontinuous steps. This enables the phases to overlap where needed allowing requirements to be gathered while overall project progress is still proceeding.

- Evolutionary Prototyping involves multiple iterations of requirements gathering. Iterations produce individual prototypes to be presented to the customer to stimulate further feedback and discussion of requirements.
- Code-and-Fix is the typical approach to avoiding the complexities of a development methodology. It is only useful for small, throw away projects and is dangerous because it offers no Quality Assurance or Risk Management.
- Staged Delivery involves breaking design, coding, testing and deployment into separate stages which are useful to the customer. Each stage must function independently of other stages.
- Evolutionary Development straddles evolutionary prototyping and staged delivery. Initial development is on lower-level functions which will hopefully remain independent of changing customer requirements. (Business ESolutions 2002)

The Modified Waterfall Model is preferred in the project and will be employed as long as the project remains on schedule. Failsafe will be provided by Feature Driven Development if the project schedule is not fulfilled.

3.3 Programming Language

Business today demands sophisticated computing capabilities. Even the software products used for office automation (word processors, spreadsheets, etc.) have become large and complex in the process of meeting user needs. The issues involved in creating large, complex software are many and varied. However one issue continues to cause controversy and seldom results in agreement between programmers. This issue is the choice of Programming Language. Some alternative programming languages are presented and discussed in this section.

3.3.1 Visual Basic.NET

VB.NET was developed as part of Microsoft's Visual Studio solution and represents the next generation of language and tools for rapidly building Microsoft Windows and Web applications. VB.NET has a very clean interface for designing a Graphical User Interface (GUI) making it extremely popular when developing such applications.

Unfortunately VB.NET is not a common tool in most Computer Science and Engineering applications. Some features are useful for special Computer Science applications, however it is more generally regarded as a business programming language. For this reason VB.NET was disregarded despite supporting the necessary socket operations required for this development project.

3.3.2 ASP.NET

ASP.NET is derived from the Active Server Pages language used to create dynamic web pages. This is not exactly suitable for a Network Address Translator which normally works at a much lower layer than ASP.NET. However ASP.NET would serve a useful function for this project as a web interface for management of the Network Address Translator (NAT). A particularly useful aspect of ASP.NET would be that a web interface could be used to allow an Application Layer Gateway (ALG) to report an anticipated incoming port. This may not be useful in all cases because the requested port may already be in use. However it may be used to attempt a repair of some applications which do not normally work under NAT such as Active FTP.

ASP.NET will not form part of this project however because it is unlikely to reach the application layer. The most important issue is to achieve a simple working NAT for use in future development.

3.3.3 C

C evolved from a language called B, written by Ken Thompson. C is a simple and small language, which can be translated with simple, small compilers. Today it is among the

languages most commonly used throughout the computer industry.

There is no particular reason why this project could not have been developed in C. It has a clean interface, is easy to use and supports all the necessary sockets operations. The only downside is the difficulty in generating a Graphical User Interface (GUI) using appropriate libraries. Writing a windows service in C can be difficult as a number of hooks need to be developed. These hooks are for use by the operating system in starting and stopping the service.

Although C was not used in this project it was excluded only because of time considerations in GUI development and some difficulties in writing a Windows Service. The final product could be ported back to C as a future project.

3.3.4 C++/Visual C++.NET

C++ is a rewritten and improved version of C. The major focus in developing C++ was to enable the development of object oriented programs. C++ also showcases a variety of other features not found in C while still maintaining the basic syntax and semantics found in C. Visual C++.NET is a particular implementation of C++ by Microsoft. Visual C++.NET includes integrated GUI development tools which make developing a GUI type interface much easier than using standard libraries in C or other C++ environments.

Again there is no specific reason why C++ could not be used to develop this project. The added benefit of integrated GUI development tools in Visual C++.NET only adds to the reasons for using C++. However Visual C++ still requires complex methods to support running as a Windows Service. In addition the Microsoft Foundation Class (MFC) is extremely difficult to master and takes a lot of time to set up properly.

Future work could involve porting the application back to Visual C++ or developing a user interface using standard C++ libraries.

3.3.5 Java

Java is an object oriented programming language which uses a Java Virtual Machine (JVM) to run Java programs. The JVM is cross-platform capable having been ported to many variants of Windows and Unix. The Java Application Programming Interfaces are a set of pre-built classes that can be used in program development.

In C++, memory has to be explicitly requested when required. Likewise, when finished with the memory, it has to be explicitly returned to the operating system. Although this process sounds simple, it is easy to create a memory leak, which is when your application requests memory and forgets to release the allocated memory. Over time, the application grows in size, slows down the system greatly and eventually crashes. Java implements a feature called Garbage Collection which automatically recovers memory that can no longer be referenced by the program. The result is that memory leaks cannot occur, to an inexperienced programmer it looks like every variable is causing a memory leak, however, behind the scenes the garbage collector is searching for any piece of memory which is no longer required and reclaiming it to be reassigned to another variable or even another program. (Campione, Walrath & Huml 2000)

Despite the feature of garbage collection in Java it was not used in this project. Java only supports two types of sockets, Transmission Control Protocol (TCP) and User Datagram Protocol (UDP). A NAT requires a lower level socket known as a RAW socket which receives an entire packet with the TCP and Internet Protocol (IP) headers intact. Despite some later references to using a lower level protocol to inject packets at the data link layer the decision to reject Java as a programming language had already been made.

3.3.6 C#.NET

C#.NET, a Java-like programming language, was developed by Microsoft and submitted to the ECMA standards group for approval. Although the language is Java-like its syntax and semantics are remarkably similar to C/C++. The likeness to Java comes from many features which are implemented in Java such as garbage collection and the

implementation of a hashing-table in standard functions for almost every variable type.

C#.NET had several other advantages over most other languages. It supported creating a Windows Service using its own implementation. The only feature left for development was code to start and stop the service. All the necessary hooks were linked to the developers functions. In addition C#.NET supports a very wide range of sockets including RAW sockets which can send and receive packets including the full IP header. The GUI development platform used in VB.NET is reproduced almost perfectly in C#.NET enabling the simple development of a simple management interface for the project in the same language in which the service was developed.

C#.NET was my final language of choice for this project. Although there have been critics of C#.NET my decision was to use it in a kernel level project. The language is quickly becoming popular and has successfully been used in a large number of commercial projects.

3.4 Chapter Summary

This project followed the Evolutionary Development Model. By developing prototypes commencing at low-level functionality and increasing this functionality until the overall requirements are met, project progress can be measured through the attainment of milestones and project research can be broken into individual concepts.

C#.NET is not the most common language for low level networking applications. However it was the programming language with the most features that were of practical use in this project. Any other language would have required significantly longer periods of time for development however C#.NET was similar enough in syntax to C/C++ that the learning curve of the language was very short.

Chapter 4

Internet Protocol

4.1 Chapter Overview

No Network Address Translator (NAT) could be developed without a full understanding of the Internet Protocol (IP). Already IP has been mentioned in several sections of this book. It has been largely undefined, except to say it is a protocol developed and used in the Internet. However, the use of IP is not limited to the Internet, many smaller networks also use IP as an underlying communication protocol. This is testimony to the robustness of IP, it can be used to deliver packets to the next cubical in the office, across the street, across the country or around the world. This chapter examines the reasons for developing the Internet Protocol, why it is so useful in computer communications and the reason such a protocol would need added functionality in the form of a NAT. Finally, the report shall briefly discuss what is currently being developed to remove the need for a NAT and return IP to a completely independent protocol.

4.2 IP Functionality

The Internet Protocol is the glue that holds the Internet together. It was designed explicitly for the task of internetworking or connecting many different types of network. Its primary function is to provide a best-effort attempt to deliver segments of data

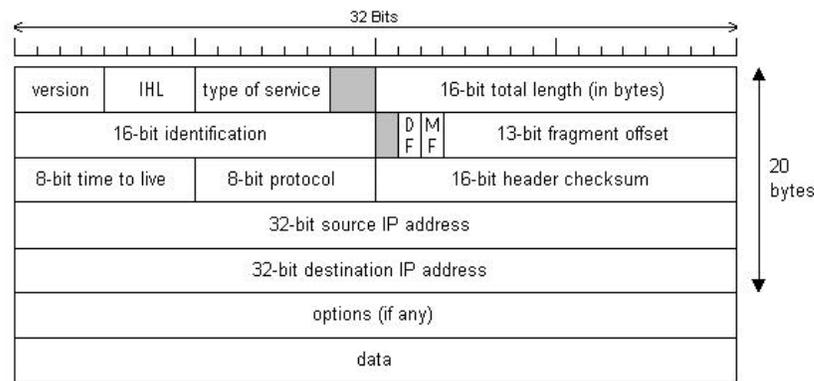


Figure 4.1: The IPv4 (Internet Protocol) header. (adapted from (Tanenbaum 2003)).

called datagrams from source to destination with no regard for where the source and destination are, be it on the same corporate Local Area Network, different Local Area Networks that are interconnected by a router or on two completely different networks separated by a multitude of other networks which are connected by many different paths, commonly called a route. (Tanenbaum 2003)

4.2.1 Looking at an IP Header

Each datagram being sent via IP has a special IP Header added. The header is used to identify the packet, set some control information, record the length of the datagram, record error checking information (error checking is only for the header itself, not for any data the header may contain) and indicate the source and destination of the datagram. The header consists of a 20-byte fixed part and a variable length optional part. There are two ways of looking at an IP Header. The first method is to review the theoretical contents of the header and learn what each part indicates. The second method is to examine a real IP packet to observe the practical application of the theory. This is important in the design and implementation of a NAT because the header will not label its contents adequately to assist the user in determining what the numbers actually mean. The structure of a IP Header is shown in Figure 4.1.

Quickly Summarizing the meaning of the individual fields in an IP Header:

- The *Version* field indicates which version of the protocol the datagram belongs

to allowing transitions to new versions to roll-out over many years.

- The *IHL* or Internet Header Length is provided to define how long the header is in 32 bit words. The default and minimum value is 5 indicating a header with no options. The maximum value is 15 indicating a header containing 40 bytes of options.
- The *Type of Service* field distinguishes between different classes of service. For example real-time voice requires fast delivery, however, this is not concerned with reliability. In fact for most voice applications reliability and error checking cause more problems than they solve. For data downloads however reliability is more important than throughput (despite what some users would suggest). The actual use of this field will not be explained here as in practice most routers ignore its contents anyway.
- The *Total Length* indicates the length of the entire datagram which may consist of a IP header, options and data. The absolute maximum Total Length is 65,535 bytes which is currently suitable for 1500 byte maximum ethernet frames, however this is not ideal for new multi-gigabit ethernet connections.
- The *Identification* field allows the destination to determine which datagram a IP fragment belongs to. All IP fragments that originate from the same IP packet have the same Identification Number.
- *DF* is a single bit which stands for "Don't Fragment" and is an instruction that the packet must not be fragmented (normally because the destination does not have a full IP Stack loaded and cannot reassemble packets). IP requires that every participating network accepts frames of 576 bytes or less.
- *MF* means "more fragments". If an IP packet is fragmented all pieces except for the last will have this bit set.
- The *Fragment offset* indicates the relative position of the current fragment in the fully assembled packet. The offset is given as a number of 8-byte fields which offset the current fragment. As 13-bits is being used, 8192 fragments can occur. This allows complete fragmentation of a 65,536 byte packet (one byte larger than allowed by the IP protocol).

- *Time to live* is a counter of packet hops. It is decremented once by each router the packet passes through. If the value reaches zero the packet is discarded and the host warned. This system prevents a routing loop from buggy router configurations from crashing several backbone routers as packets flood into the loop but never leave.
- RFC1700 was the first global definition of transport level protocols. The globally accepted list is kept at <http://www.iana.org/assignments/protocol-numbers>. The *Protocol* field may contain any number from this web page. It is used at the destination and in some other circumstances such as Network Address Translators to determine what data to expect following the IP Header and Options.
- As stated earlier IP is a best effort protocol which will attempt to deliver datagrams, however makes no guarantee that individual datagrams will not become corrupted or not reach their destination. The *Header checksum* is used to verify that the data contained in the IP header is not corrupt. Higher level protocols often also use a checksum that verifies the entire packet, including the data, has not been corrupted. The IP checksum guards against routers with bad memory modules and ensures that when IP reports the source to a higher layer the address is correct.
- The *Source Address* is a 32 bit field used to identify the sending host.
- The *Destination Address* is a 32 bit field identifying the receiving host. More information on IP Addresses will be provided in a following section.
- A large range of special features have been defined for use in the IP header through the use of the *Options* field. The current list is kept current at <http://www.iana.org/assignments/ip-parameters>. The use of this field will not be discussed and has become depreciated due to the limited size of the field and the size of the global internet. (Tanenbaum 2003)

Table 4.1 shows the output of the debugger which has captured an IP packet. It demonstrates how IP Headers are practically used, however it would be a daunting task to actually decode the meaning of a packet without the details provided by the explanation.

Table 4.1: Contents of a real IP packet.

Buffer Position	Byte Contents	Explanation
[0]	69	Version = 4, IHL = 5
[1]	0	Type of service = 0 = Normal
[2]	0	Total length = 45
[3]	45	
[4]	159	ID = 40756
[5]	52	
[6]	64	Don't Fragment
[7]	0	
[8]	128	TTL = 128
[9]	6	Protocol = 6
[10]	0	Checksum = 15
[11]	15	
[12]	192	Source = 192.168.0.3 (My Computer)
[13]	168	
[14]	0	
[15]	3	
[16]	207	Destination = 207.46.106.173
[17]	46	
[18]	106	
[19]	173	
[20]	7	The rest of the packet can just be considered as data.
[21]	105	
[22]	7	
[23]	71	
[24]	231	
[25]	219	
[26]	240	
[27]	123	
[28]	109	

Table 4.1: (continued)

[29]	112	
[30]	168	
[31]	186	
[32]	80	
[33]	24	
[34]	65	
[35]	89	
[36]	213	
[37]	88	
[38]	0	
[39]	0	
[40]	80	
[41]	78	
[42]	71	
[43]	13	
[44]	10	

4.2.2 Fragmentation

IP was designed to work over many different types of networks with various hardware because it must accommodate for differences in the maximum frame sizes due to different underlying networks. The maximum frame size of the underlying network topology is called the maximum transmission unit (MTU).

Suppose in Figure 4.2 that Host A wants to send a large amount of data to Host C, hence the IP Protocol creates a packet which is 1500 bytes long (1480 bytes of data + 20 bytes for the IP Header). The router at Host B receives the packet and assuming the Don't Fragment bit is not set, it will create two packets both destined for Host C. The first packet will contain 976 bytes of data (the maximum multiple of 8 bytes + 20 bytes for the IP Header that can fit on the second network), will have the More

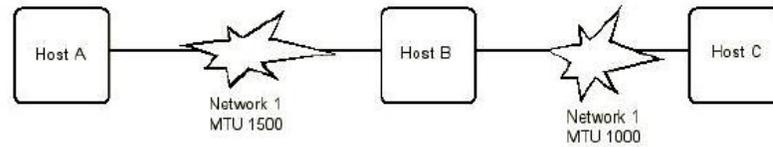


Figure 4.2: An Example of connecting two networks with differing MTU values. (adapted from (Feit 1998)).

Fragments (MF) bit set and will have a *fragment offset* of zero. The second packet will contain the remaining 504 bytes of data, will not have the More Fragments (MF) bit set however will have a *fragment offset* of 976. Host C will then have sufficient information to reassemble both fragments and receive the correct data. (Tanenbaum 2003)

4.2.3 IP Addresses

An IP Address is a 32 bit number used in the source and destination fields of an IP header. Each IP address consists of a network number and host number. In principle no two machines on the internet can have the same IP address (in practice this situation could occur however one host will not receive the packets unless the routing tables are configured incorrectly). It is worth mentioning that one of my home computers is a router and therefore receives two globally unique IP addresses, one for each of the external interface to the internet and internal interface for the intranet (internal network). When these issues are considered, the diminishing number of unused IP addresses is not surprising.

IP Addresses are usually written in dotted decimal notation. For example the 32-bit Hexadecimal address C0A80003 (or in decimal 3232235523) is written as 192.168.0.3. Several IP addresses also have special meanings. For example the IP address 0.0.0.0 refers to the current host. In the above mentioned network the address 0.0.x.x refers to the host with the given IP address on the local network. The address 255.255.255.255 indicates the broadcast address which means all hosts on the local network. The broadcast address on a remote network in the same class as above would be x.x.255.255 (though most administrators disable such addresses because they are a security risk).

The IP address 127.x.x.x always refers to the loopback device. The loopback device is a method of sending packets to the local machine without putting the packet onto the physical wire. The loopback address is also valid for a device which may not have any network interface installed to ease testing requirements. (Feit 1998)

4.2.4 Subnets

As mentioned earlier IP addresses encode the network and host number. For example the 192.168.0.3 IP address given consists of the network address, 192.168 and the host address, 0.3. Any computer in the same network must also have the same network address with a different host address. Varying each part of the host address from 0 to 255 gives 65025 addresses. Even though some addresses are reserved for the special purposes there are over 60000 usable IP addresses in the given address range. Unfortunately Ethernet was designed with much stricter limits of only 1024 hosts per network. The problem is that the networking authority will not give out two networks of 65025 hosts each simply because the underlying network was not scalable (especially when IP Addresses are already becoming scarce).

The solution was to allow networks to be split into several sub-networks or subnets although still appear as a single large network to the wider Internet. In the above example a subnet address would be specified to segregate the larger network into several smaller networks. For example a subnet address of 255.255.255.0 would allow 256 subnetworks (in some cases subnet masks of all zeros or all ones in the address cannot be used reducing this value to 254) each containing up to 254 hosts (0 and 255 are reserved in each subnetwork as mentioned previously). (Feit 1998)

4.2.5 Classless InterDomain Routing

Classless InterDomain Routing (CIDR) is a solution analogous to scraping the very last soup from the bottom of the pan. CIDR suggests allocating remaining IP addresses in variable sized blocks. The blocks are not completely free from restrictions, they still need to be allocated in blocks of 2^x however this allows the last of the IP address

space to be allocated based on the current proven needs of an organization rather than anticipated future needs. Research has shown that over 50% of all networks supporting 64K hosts actually have less than 50 hosts. Under the old method of IP Address allocation these networks could have been given a network address supporting only 254 hosts and wasted IP Address space would have been avoided.

With CIDR each routing table entry has a 32-bit mask added. When a packet is received by the router the destination address is extracted and each address in the routing table is compared to the destination address which is masked by the correct mask (one bits in the mask allow the equivalent bit in the destination address while zero bits hide it) from the routing table. When a match is found the correct forwarding interface is looked up in the table and the packet transmitted to the next hop. Because multiple entries may match the destination address due to different masks the longest masks are used first. A long mask (a large number of the 32 bits are ones) indicate a very specific network whereas shorter masks indicate a more general case.

Fortunately the routing does get easier after the specific cases. Special cases are normally caused when the packet is very close to its destination. Consider the extreme example when the next hop will be to the final host, the router must decide which of several ports leads to the destination and will forward the packet on that port only. Now consider the other case where the destination is a long way from the current location (in terms of number of hops). When this occurs the routers may generalize specific cases. For example imagine that all internet addresses starting with 203.2.x.x belong to Australians. A router in Los Angeles does not need entries for 203.2.1.1 and then 203.2.1.2 when it can have a single entry for 203.2.0.0 with a mask of 255.255.0.0. All packets matching this routing entry would likely be forwarded by a transpacific ocean cable which might be called pacf01 by the router. (Tanenbaum 2003)

4.3 Network Address Translation

Despite the efforts to stop wasting IP Address space using solutions such as CIDR and requiring organizations to prove the need for new IP Addresses there has been

little to no attempts to resolve the problem. New Internet designs with much larger IP Addresses are being tested, however final commercial implementation will be years away, if ever. A quick fix is required that can be implemented anywhere across the Internet at any moment. This quick fix has been developed in the form of Network Address Translation (NAT).

4.3.1 Overview of NAPT

Two types of NAT have been developed. The first version, Traditional NAT does not save IP addresses for end users. It can however be implemented by an ISP to save on the IP addresses wasted by allowing many small CIDR subnets (remember that a subnet with 2 IP addresses actually consumes 4 IP Addresses). Traditional NAT simply translates the users IP address to a globally unique IP address for packets destined for the internet. Hence if a user needs 4 IP addresses they are allocated 4 private IP addresses each being mapped to a global IP address by the service provider.

The second method of NAT is more accurately called Network Address Port Translation (NAPT). NAPT is concerned with actually reducing the need for multiple IP addresses. The idea is to multiplex a number of access requests to the external network onto a single globally unique IP Address (although there is no rule stating that only one global IP address could be used for NAPT). The incoming packets to the NAPT are then demultiplexed back to the original source. Internally this is accomplished by mapping tuples of the type (local IP addresses, local TU port number) to tuples of the type (registered IP address, assigned TU port number) where TU is the transport layer unit. Supported TUs are normally Transmission Control Protocol (see 5), User Datagram Protocol and Internet Control Message Protocol query's only. Limited inbound access can be provided by statically mapping a known TU port service to a specific local IP address. (Srisuresh & Egevang 2001)

4.3.2 Address Binding

In NAT implementations, binding would take place between the tuple of (private address, private TU port) and the tuple of (assigned address, assigned TU port). This binding is created when the outgoing session commences.

4.3.3 Address Unbinding

NAT may unbind the tuple of (assigned address, assigned TU port) when the last connection closes. However because the situation of a host crashing must be handled a timer is required to unbind addresses after a period of inactivity.

4.3.4 Header Manipulation

In the IP layer every packet header must be modified. These modifications are for the source IP address for outbound packets and destination address for inbound packets. The IP checksum must also be updated.

For TCP/UDP protocols the source port must be updated for outbound packets and restored for the destination port of inbound packets. The checksum must also be updated remembering that a pseudo header including the IP addresses forms part of the checksum. As an exception UDP packets with zero checksum should not be updated.

ICMP packets must also be specifically updated for the purposes of NAT. The required updates are to the ICMP Query ID and ICMP checksum. The Query ID must be modified from an internal ID to assigned ID for outbound packets and assigned ID to internal ID for inbound packets. (Srisuresh & Egevang 2001)

Sometimes it may be necessary to add a specific protocol to a NAT. The implementation details are not defined by the NAT standards and will vary depending on the requirements of the protocol and the best practices for translating the protocol at the time the supporting software is written.

4.3.5 Incremental Checksum Adjustment

IP, TCP, UDP and ICMP headers all use the same form of checksum. Unfortunately the calculation of these checksums from scratch is computationally expensive. Fortunately there is a better method of calculating an incremental checksum for an existing header. There are two advantages of using this incremental checksum. First updating the checksum based on modifications to the header is less computationally expensive than recalculating from scratch. Second this form of update avoids the need to check for packet corruption. If the received packet is corrupt and the checksum is recalculated the next router or the destination machine will incorrectly believe the packet is correct because the recalculated checksum will equal the true checksum of the now corrupt packet. However incremental checksum on a corrupt packet will not correct for the corruption so the next router or destination that attempts to verify the checksum will fail and drop the packet. (Rijsinghani 1994)

4.3.6 ICMP error packet modifications

There is a slight difficulty in NAT caused by ICMP error packets. The ICMP error packet may contain an embedded IP packet (normally one which caused an error) and this IP packet will contain source and destination addresses and possibly TCP/UDP port numbers which need to be updated to ensure end-to-end transparency of the NAT system. In addition the checksums of all embedded packets must also be updated to reflect changes made by modification of the ICMP error packet. More details on ICMP packets can be found at <ftp://ftp.rfc-editor.org/in-notes/rfc792.txt>.

4.3.7 FTP Support

NAT requires special consideration when using File Transfer Protocol (FTP) in combination with NAT. This is because FTP encodes several pieces of IP data into control packets. This data can include IP Addresses and TCP Port numbers. Additionally FTP can negotiate to open another port for the following file transfer. NAT needs to recognize this expected incoming port or packets will be lost as they are destined to a

(assigned address, assigned TU port) tuple that NAPT did not assign.

More details on the exact requirements of an Application Layer Gateway to overcome this problem are provided in Srisuresh & Egevang (2001).

4.3.8 Using IP Options

Although IP Options are depreciated and seldom used there is a possibility that private addresses contained in an IP Option would remain untranslated in a NAT packet traversing the Internet. This problem is not addressed in many implementations of NAT because routers using IP options should only consider the next-hop and the presence of a private IP address would be overlooked.

4.3.9 Recommendations for Private IP Address Range

Organizations using NAT are recommended to make use of the three private IP Address ranges provided by the Internet Assigned Numbers Authority (IANA). These Addresses are 10.0.0.0/8, 172.16.0.0/12 and 192.168.0.0/16. If these numbers are not used one possible problem that can occur is the local machine completes a DNS lookup and finds a local IP address is the destination. However the NAPT will be unable to differentiate between a request to the local IP address and one that should have been translated to the Internet.

4.3.10 Privacy and Security

NAPT provides a privacy and security mechanism by shielding internal clients from any unexpected inbound packets. Inbound packets are only possible if a port in use by NAPT is selected for the attack and the attack source uses the same IP address as the external machine. Launching such an attack is extremely difficult because both details are normally stored on the NAPT servers memory and are not available to users or external computers. Also because use of an invalid IP address is normally required by the attacking machine two way communication cannot normally be established.

However NAT does have the undesirable impact on internal policing. If an internal client uses NAT to shield their attack on the Internet the owner of the NAT server will normally be blamed for the attack. Unless detailed logs are available the offending person cannot usually be identified.

4.3.11 Fragmented Packets

NAT will never be able to successfully translate outbound TCP/UDP fragments. This failure results in the TCP/UDP header being contained in one of the fragments and not in any other fragments. Normally the IP Fragment number would be used in this situation however there is no guarantee that two client machines will not use the same fragment identifiers and result in corruption. The only solution is to have NAT reassemble fragmented IP packets before allowing translation. This solution is also recommended for the purposes of enabling Secure IP (IPSec) over NAT. (Srisuresh & Egevang 2001)

4.4 IPv6

One of the newest major standards on the horizon is IPv6. Although IPv6 has not officially become a standard, it is worth some overview, especially since the final introduction of IPv6 will likely make the outcomes of this project worthless. It is very possible that this information will change as IPv6 moves closer to standardization, so this is a guide into IPv6, not definitive information. (Tanenbaum 2003)

Some of the benefits of IPv6 include greater addressing space, built-in Quality of Service (QoS), and better routing performance and services. However, a number of barriers must be overcome before the implementation of IPv6. The biggest will be what the business need is for moving from current IPv4 to IPv6. The killer application for IPv4 has not appeared yet, it may not appear at all. However IPv6 will gain momentum quickly if such an application is developed. The total lack of IP addresses may eventually force the role over without commercial support. Companies will follow more out of need for compatibility than some great new web application.

4.5 Chapter Summary

In summary IP is a protocol developed for use in interconnected packet-switched networks. It provides the underlying structure necessary to obtain information from one host to another while dealing with anything that might separate the source and destination hosts.

This chapter examined the IP Header and its data contents including some of the meaning of the data. Also some of the special functions of IP to deal with common situations with internetworking. An overview of some of the issues in distributing IP addresses to various organizations was given. The issues of Network Address Translation were then presented and reviewed in detail. This section is particularly complex and borrows some details from the next chapter on the Transmission Control Protocol (TCP). Readers who found this section difficult should not be disheartened as it requires a great deal of technical knowledge. It is recommended that the chapter on TCP is reviewed before returning to NAT.

Readers who have followed the contents of the current chapter will find the next two chapters on the Transmission Control Protocol and some existing implementations of NAT much lighter reading.

Chapter 5

Transmission Control Protocol

5.1 Chapter Overview

Transmission Control Protocol (TCP) builds on the IP layer to provide end-to-end connectivity for packet-switched networks. TCP must compensate for the lack of reliability in the IP layer and must operate in such a fashion as not to overload the underlying protocols or core network routers.

This chapter will review the major features of TCP including the header structure. The concepts of connections and establishing connections will be examined and a brief overview of another Transport layer protocol called User Datagram Protocol.

5.2 TCP Functionality

TCP is a connection-oriented reliable service designed as part of the Transport layer. TCP is connection oriented, with each application employing a TCP connection between itself and the opposing TCP end-point. This must occur before data may be sent or received. The common term for this type of connection is a state based service because the state of the connection is maintained and determines the ability of applications to use the service.

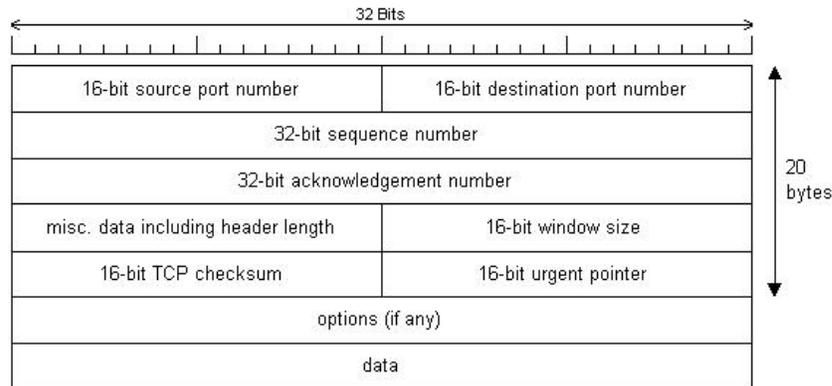


Figure 5.1: The IPv4 (Internet Protocol) header. (adapted from (Tanenbaum 2003)).

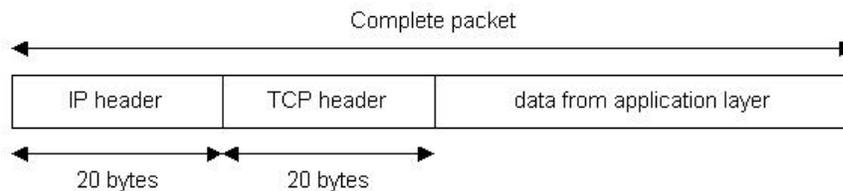


Figure 5.2: Combining TCP and IP to encapsulate data.

Figure 5.1 shows the structure of a TCP header while Figure 5.2 shows how a TCP header is added to the data which is then passed to the IP layer where an IP header is added.

Quickly Summarizing the meaning of the individual fields in an TCP Header:

- The *Source Port* identifies the end-point of the connection. The end-point must be identified to enable the TCP layer to determine which upper layer service requiring the data. Common port numbers are defined at <http://www.iana.org/assignments/port-numbers>.
- The *Destination Port* identifies the target end-point. A host's IP address and TCP Port together form a 48-bit unique end-point. Together the unique source and destination end-points identify the connection.
- The *Sequence Number* identifies the TCP segment to enable packets to be re-assembled to form a reliable byte stream.

- *Acknowledgement Number* is used to specify the *Sequence Number* of the next TCP segment expected (not the last segment that was correctly received).
- The misc data field starts off with a 4-bit *TCP Header Length*. This is followed by 6-bits which are not used and must always be zero. Finally there are six 1-bit flags used to signal important events to the TCP receiver. In order these are the *Urgent flag* (URG) which indicates that the Urgent Pointer is in use. The *Acknowledgement flag* (ACK) is set if the Acknowledgment Number is valid or cleared if the packet does not contain an Acknowledgement. The *Push flag* (PSH) is used to push the packet through the TCP layer immediately even if this means avoiding buffering. The *Reset flag* (RST) is used to reset a connection for which state information is invalid. This may be due to a host crash, to reject invalid segments or refuse a connection request. The *Synchronize flag* (SYN) is used to establish a connection. If the SYN bit is set and the ACK bit is cleared the request is to establish a connection. If both bits are set the packet represents acceptance of the connection. Finally the *Finish flag* (FIN) terminates a connection. However after sending a FIN packet the TCP port must remain open to incoming data indefinitely (assuming data continues to arrive).
- The *Window Size* is used for flow control as an indicator of how many bytes may be transmitted starting from the acknowledged segment.
- A TCP *Checksum* provides extra reliability. The TCP header, data and pseudo header participate in the checksum calculation. The checksum must be calculated with the checksum field set to zero. The pseudo header contains the IP source and destination addresses, a byte of 0 bits, a byte containing the protocol number (for TCP the protocol this is the binary representation of 6) and the TCP segment length (the IP total length minus the IP header length). Using the pseudo header violates the independence between the IP and TCP layers, however it offers additional protection against misdelivery of packets. The same concept of a pseudo header is used in the User Datagram Protocol (UDP).
- The *Urgent Pointer* indicates the byte offset from the current sequence number to find what is marked as urgent data. Urgent data is similar to computer interrupts, allowing the sender to send a type of interrupt signal to the receiver without TCP

being aware of the reason for the urgent data.

- The *Options* field in TCP offers similarity to the IP Options field which allowed the implementation of a large range of special features. However unlike IP the TCP *Options* field is extensively used by features that support newer modern networking hardware and software. (Tanenbaum 2003)

5.2.1 Addressing

TCP provides its own addressing mechanism. Unlike IP however TCP addresses are not normally used by intermediate routers and are only meaningful to the end-host. Addressing is required to identify the application or service to which the data belongs. (Day & Zimmermann 1983) Some applications have default destination addresses such as HTTP. When requesting a web page we do not need to specify the use of port 80. However `http://www.google.com.au:80` is a valid reference to the Google website. With 65025 possible port numbers an internet host would need to run several thousand simultaneous web applications to use all possible ports. This knowledge allows the implementation of a Network Address Port Translator which uses ports to identify different hosts, a job normally left to the IP layer.

5.2.2 Reliability

TCP offers a reliable end-to-end byte stream which means each byte must be delivered once, in order and without error. However the IP layer does not guarantee any of these features. IP packets may be duplicated, lost, corrupted or delivered out of order. The TCP protocol uses sequence numbers, checksums that include the data and acknowledgement to ensure that received data is correct. Any TCP segment received with an incorrect checksum is immediately discarded. Such segments cannot be requested again at the time they are discarded as the checksum error could be due to corruption of the port numbers or sequence number. Sequence Numbers are also used to ensure correct ordering of the data and for the purposes of Acknowledgements. Acknowledgements may only be sent for segments which have been successfully received in order, TCP may internally buffer segments which are separated from the current sequence

by a gap however may not transmit acknowledgments until the gap has been correctly filled. TCP may request that segments causing gaps in the stream are resent explicitly or request that all segments following a specific sequence number are retransmitted. However the sending host is responsible for retransmitting packets based on a timer. Support of requests from the opposing TCP end-point is optional. (Tanenbaum 2003)

5.2.3 Congestion Control

Internally the TCP recognises two types of congestion. Network Capacity is a response to an overflow of packets on the network which will cause routers to run out of buffer space and internally discard packets before they are delivered to the destination. Receiver Capacity is the amount of data the opposing TCP end-point is prepared to receive before it requires time to process the received data and pass it to the upper layers.

Internally TCP maintains two windows. One is the amount of data the receiver is prepared to accept, the second is the congestion window. TCP will never send more data than is indicated by the minimum of these two windows. The receiver window is controlled by the remote end-point. However controlling congestion on the network is slightly more involved.

Congestion control uses the congestion window and a threshold to find an optimal amount of data to fill the network without causing congestion. The threshold is initially set at 64 KB and TCP is allowed to transmit 1 KB of data. The TCP retransmission timer is started and TCP waits for acknowledgements or a timeout. Whenever the amount of data sent is all acknowledged without an error or timeout, TCP is allowed to transmit double the amount of data as on the previous attempt unless the amount has reached the threshold. After the amount of data sent on the previous attempt is equal to the threshold both grow linearly at a much slower rate (about 2 KB for each successful transmission). It is important to note that the amount of data transmitted may never exceed the threshold as the threshold grows at the same linear rate. When a timeout occurs the threshold is halved and the process repeats with TCP being allowed to send 1 KB which is doubled on every successful attempt until the threshold

is reached. The process repeats indefinitely ideally with an average of about the correct amount of data the Internet can correctly handel at the time. At all times TCP may never send a burst larger than the minimum of the congestion control size and the receiver window size. (Tanenbaum 2003)

5.2.4 Connection Management

TCP needs to initialise and maintain some state based information including local and remote end-point information, local and remote sequence numbers and window sizes. Each TCP connection must be initiated and initialized by the transfer of state information. For example TCP hosts must agree on window sizes and starting sequence numbers when the connection is created. When the connection is closed the memory used for this state information may be freed for other applications or connections. All connection management data transfer is also checked and TCP must recover from any error such as lost connection requests in a timely manner. (Robison 2002)

5.3 Data Connections

5.3.1 Three Way Handshaking

All TCP connections begin passively with one side waiting for a connection. This involves a blocking call to the LISTEN or ACCEPT methods specifying a particular source address or accepting connections from any source address. The call is named as blocking because the application cannot proceed until a connection is made, the execution of the code is blocked until a connection request is received.

The second step is for an application to execute the CONNECT method specifying the end-point of the connection (IP Address and TCP Port), the window size and optionally any user data to be used when establishing the connection (a username and password for authentication perhaps). TCP then sends this information out in a packet with the SYN flag set. The recieving machine checks that an application has executed a LISTEN or ACCEPT method on the specified port. If this does not occur, the connection is

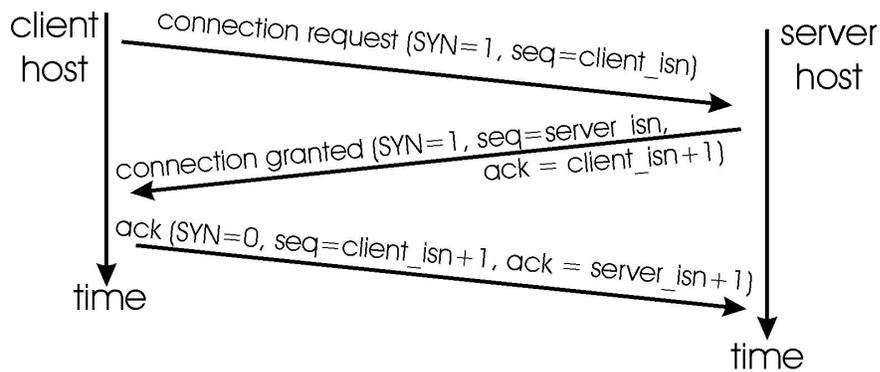


Figure 5.3: TCP connection establishment.

rejected by replying with a packet that has the RST flag set. Otherwise, the application that executed the LISTEN or ACCEPT method receives the connection request. This is accomplished by sending a reply packet with both the SYN and ACK flags set and the acknowledgement number being the sequence number used in the connection request plus one to indicate the next segment. Finally the session is fully opened by a reply with the next sequence number and the ACK flag set for the response. (Tanenbaum 2003) This process is depicted in Figure 5.3.

5.3.2 Simultaneous Open

In TCP there is the possibility that two hosts will simultaneously attempt to establish a connection. If this occurs only one connection must result. TCP handles this situation by ordering the connection end-points. Hence both connections will result in the connection (x,y) never (y,x) and TCP will only record one table entry for the connection. Each host will reply by resending their initial SYN while including an ACK for the opposite host's request. This situation is shown in Figure 5.4.

5.3.3 Active Close

TCP Connections may be closed by two methods. The first is an active close where the client closes the TCP application causing a FIN segment to be sent. The client may then receive a FIN only segment indicating a simultaneous close which is similar

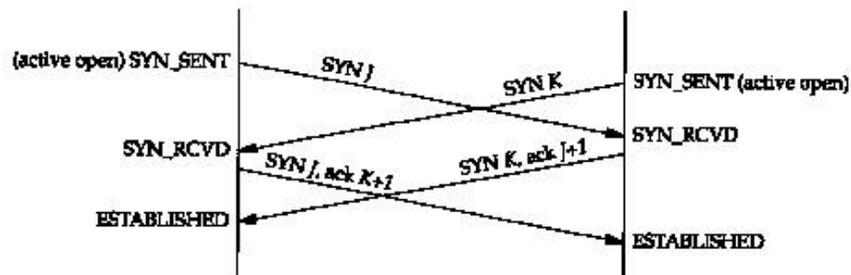


Figure 5.4: TCP Simultaneous Open. (adapted from (Tanenbaum 2003))

to a simultaneous open. Both applications then send ACK segments to finalize the connection. Alternatively the client may receive a FIN packet also Acknowledging its FIN. The client will send its final Acknowledgement and close the connection. Finally the client may receive an Acknowledgement to its FIN without a FIN from the opposing end-point. This indicates that the opposite end-point may have more data or has kept the connection open for some other reason. The client must keep to its promise that it has completed sending data, however must continue to accept data until the opposite end-point agrees that it is also Finished (FIN) at which point the connection may be closed. In all cases a closed connection must be retained in memory in case any lingering packets related to the connection arrive.

5.3.4 Passive Close

Passive Close is the alternative method of ending a connection. A Passive Close occurs when the host receives a FIN before the application terminates. The host replies with an ACK segment for the FIN, however it must then wait for the application to complete its communications which may include sending more data before the return FIN can be transmitted. Once this FIN is Acknowledged the connection may be released.

5.4 User Datagram Protocol / Real-Time Transport Protocol

The User Datagram Protocol offers the benefits of IP communication without the strict connection establishment, state and release requirements involved in TCP. The main reason to offer this service is the addition of the port information found in the UDP header. Without this port information the transport layer would not be able to determine the appropriate higher layer destination of the information.

UDP segments can choose to use a checksum to ensure reliability. This option can also be disabled by setting the checksum to zero. Sometimes using the checksum is of little use due to the fact that the data stream is real-time and although the processing algorithms can deal with corrupt data it cannot wait for the data to be re-transmitted.

Another important feature of UDP is the lack of flow control, error control or retransmission. Particularly the lack of flow control means data is pumped onto the network as fast as the network layer can manage. This is particularly important for real-time communication.

The advantages of UDP in real-time applications lead to the development of an additional underlying protocol called the Real-Time Transport Protocol (RTP). The basic function of RTP is to multiplex several real-time data streams onto a single outgoing stream. RTP works on top of UDP which does raise some questions regarding as to which layer of the OSI or TCP/IP Reference Models it belongs.

RTP supports features such as time-stamping, encoding identification and sequencing. Sequencing is only concerned with the correct order of packets and helps the target applications account for lost data. For example if a video frame is lost it might be better not to update the video feed until the next frame arrives as opposed to blanking it. However, if the video is encoded, special consideration to prevent the corruption spreading across future frames may be required. There is also a Real-Time Transport Control Protocol which handles feedback from the RTP protocol. (Perkins 2002)

5.5 Chapter Summary

This chapter looked at Transport Layer protocols, an important link in computer communication. TCP provides a full duplex, reliable, flow controlled service to higher layers. TCP establishes and maintains connections on behalf of the higher layer service and then allows the transfer of data with the remote end-point. The data is delivered at the remote end-point and organized into correct order with any duplicates removed. TCP also handles connection release when both applications have finished using the communications channel.

TCP also plays a important role in the successful implementation of a Network Address Translator (NAT). By recognizing that TCP port numbers can be used by the NAT to represent internal end-points we can identify several different connections by varying the port number used for communication and translating this port number to the correct IP Address and Port Number at the NAT boarder.

It is important to recognize that several other transport protocols exist such as UDP. For a complete Network Address Translator we need to make as many of these protocols as possible function correctly under NAT. This is especially important for protocols which are in common use over the Internet so that NAT can provide functionality which approximates true IP Routing as closely as possible.

Chapter 6

Existing Network Address Translators

6.1 Chapter Overview

Part of implementing a new and successful Network Address Translator (NAT) is to understand current NAT implementations, the reason they are successful, in what manner they could be improved to become widely adopted and the purchase price. If a NAT can be developed to include better features, not require additional improvement and/or is cheap to customers, it will almost certainly gain a share of the market. If several of these objectives can be achieved, the current market leaders sales' would diminish as the new product prospered.

Obviously it is not the objective of this project to create a NAT which is ready for commercial distribution, however the issue of possible future commercialization of the product must be addressed for the lifetime of the code to extend over several years. If this issue is not considered the outcome of this project may not be suitable for commercial production environments and the entire project will become a throw-away prototype.

6.2 Windows

6.2.1 NAT32E

NAT32E is an enhanced IP Router allowing all private hosts on one or more Local Area Networks (LAN) to access the internet. NAT32E supports a range of connection interfaces including Dial Up Networking (DUN), Cable Modems, Asynchronous Digital Subscriber Line (ADSL) interfaces or Remote Access Service (RAS) interfaces. Configuration is automatic on most systems while manual web based configuration is also supported.

NAT32E supports a new feature called Connection Aggregation. This allows the NAT server to split data requests among two or more dial up modems. This feature may appear very useful, however it does not provide any level of competition to new Broadband services which operate without exclusive use of a phone line and cost approximately 50% less for the same speed.

NAT32E has many of the features that make NAT software most popular with home and small office users. It retails for US\$50 for the more advanced version and US\$25 for the single network only version.(NAT Software 2004)

6.2.2 BrowseGate 3 NAT/Proxy server and firewall

Browsegate 3 provides easy to use access to the Internet for all networked PC's. This includes all common services such as Web, Post Office Protocol (POP) and Simple Network Mail Protocol (SNMP) e-mail, Network News Transfer Protocol (NNTP), File Transfer Protocol (FTP) downloads or uploads and streaming video, audio and chat programs.

Browsegate 3 includes an integrated Firewall to stealth selected inbound or outbound access on specific ports. This type of technology is popular for larger organizations because it provides a higher measure of security. By controlling which services clients can access at a single point (the Firewall) the organization can make user-wide policy

changes through a single update to the settings.

BrowseGate 3 is more commonly associated with large business as is indicated by its pricing schedule and enhanced security features. Pricing ranges from US\$114.95 for a 5 computer licence to US\$1100 for a unlimited computer version.(NetcPlus Internet Solutions 2004)

6.3 Linux

6.3.1 IP Masquerading

IP Masquerading is a form of Network Address Translation developed for Linux. The goal of the package is to provide the features of high priced routers and NAT servers without the high cost. IP Masquerading maps packets from the company intranet to the Internet and maps the responses from the Internet to the company intranet.

IP Masquerading has been developed over several years and is fairly secure and stable. It is currently being used with excellent results and any new bugs are quickly fixed by the Linux development community.

Because IP Masquerading forms part of the Linux kernel it is distributed in a number of flavors of Linux which can be downloaded freely from the Internet or purchased at minimal cost in a boxed set.

6.3.2 IP Tables

IP Tables is really just a newer version of IP Masquerading used in Linux Kernels 2.4.x and above. The purpose was to create an integrated NAT and Firewall environment including the ability to forward inbound services on static ports all as part of one large system configuration. In previous versions of IP Masquerading the NAT was generally independent from the Firewall which was independent from Port Forwarding.

IP Tables used a large part of the stable and secure IP Masquerading system which

has resulted in few bugs or errors. It is more popular among high-end System Administrators who find the integrated Firewall, NAT and Port Forwarding an advantage in easing System Administration burdens.

Again this system forms part of the newer Linux kernels and is distributed freely across the Internet.

6.4 Chapter Summary

Although there are a number of free, well developed NAT Servers available, a large majority of these support only Linux based Operating Systems. Most NAT Servers developed for Windows have a financial cost associated which may vary depending on the number of machines requiring simultaneous access to the Internet through the NAT.

Most NAT Servers for windows have some features in common. They are easy to set up with few user configured options. This can be achieved in most cases by guessing which interface is to a private network and which interface is to the Internet Provider. Most feature automatic setup of clients through the use of Dynamic Host Configuration Protocol (DHCP) to assign private IP addresses to internal clients and provide information to help clients access the NAT and other services such as Domain Name Service (DNS).

The key to success of a new NAT is to ask the question, “What services could be provided by NAT which are not currently supported by existing clients?” Most NAT Servers offer a wide range of services, so what more features can the user desire? The answer to this question is largely subjective and can be broken down into two sections. Features that Business Users seek and features that Home Users desire.

Security is the most important feature to Business Users. They require internal clients to be secure from the dangers of the Internet without constant monitoring of every computer. Business Users are therefore interested in services such as Authentication (Client with IP address X cannot use the NAT until user of client X identifies himself as a valid member of the business or organization). Sometimes this leads to the second requirement of exclusion. (If user of client X has made extensive use of Internet resources then prevent user of client X from accessing the Internet for the remainder of the day.)

Home Users favour ease of use and wide application support. Application Layer Gateways for non-NAT compatible applications and special modules, to identify specific modifications required to make packets compatible with NAT, are popular among such users. Any attempts to automatically configure the NAT to support a specific appli-

cation will be well regarded by home users. High autonomy of security features and configuration will also be a popular program feature.

Obviously the requirements of Home and Business Users are very different. They may appreciate separate applications or at least different configurations based on a selection during installation. However, by tailoring the features of a NAT to more groups, increased sales and market share can be achieved. This will allow for lower cost software to compete with alternative NAT's to be viable.

Chapter 7

Network Address Translator Implementation

7.1 Chapter Overview

Implementation of the Network Address Translator (NAT) required knowledge of the C#.NET programming language and a knowledge of the Windows Socket Application Programming Interface (API) including special RAW Sockets to enable to inclusion of TCP/IP headers in the packets being sent and received. Additionally, knowledge of the function of a Network Address Translator in relation to the IP and TCP protocols is also necessary. This information concerning NAT was covered in Chapter 4.

This chapter will cover other aspects of implementation such as using the C#.NET language and Windows API Sockets.

7.2 C#.NET Basics

The C#.NET programming language is similar to an amalgamation of the C/C++ Programming Language and the Java programming language. The main difference in C#.NET is the input/output. In most C/C++ programs input/output is from the

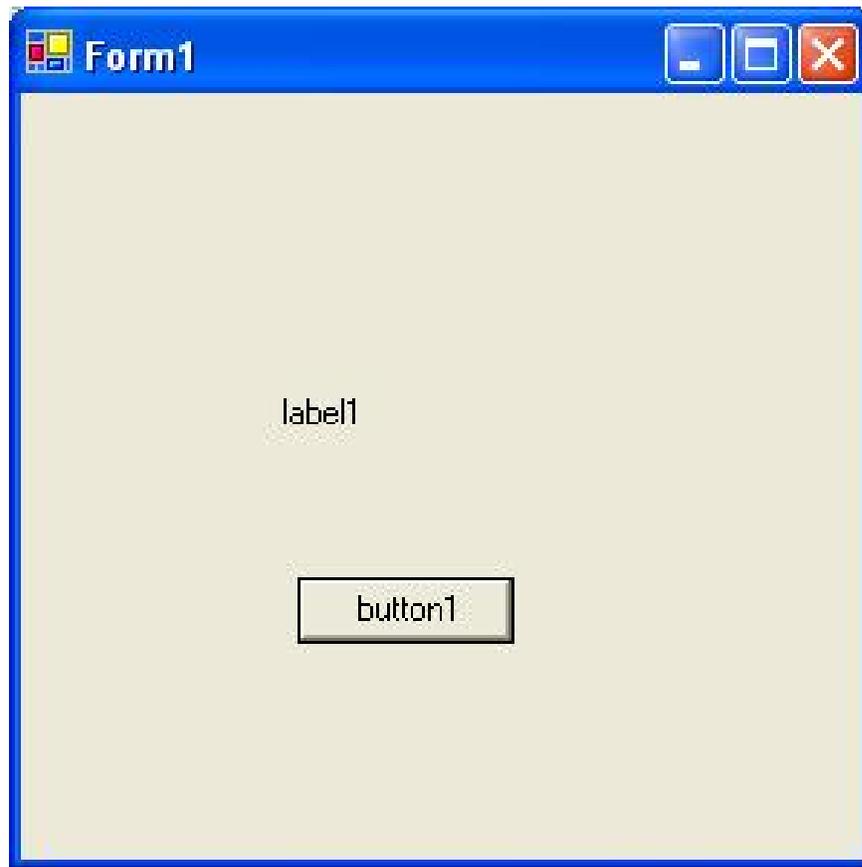


Figure 7.1: A simple C#.NET form.

console. In normal C#.NET programs there is a Graphical User Interface (GUI) and as such most input/output is from/to GUI controls. For example suppose a C#.NET Form is created as in Figure 7.1. Button1 is a input control which may be linked to some code. If we linked this button to the code in Listing 7.1 the result shown in Figure 7.2 will be displayed.

Listing 7.1: Hello World Code

```
using System;
using System.Drawing;
using System.Collections;
using System.ComponentModel;
using System.Windows.Forms;
using System.Data;

namespace Dissertation
{
    /// <summary>
    /// Summary description for Form1.
    /// </summary>
    public class Form1 : System.Windows.Forms.Form
    {
        private System.Windows.Forms.Label label1;
        private System.Windows.Forms.Button button1;
```

```

/// <summary>
/// Required designer variable.
/// </summary>
private System.ComponentModel.Container
    components = null;

public Form1()
{
    ///
    /// Required for Windows Form Designer
    /// support
    InitializeComponent();

    ///
    /// TODO: Add any constructor code
    /// after InitializeComponent call
    ///
}

/// <summary>
/// Clean up any resources being used.
/// </summary>
protected override void Dispose( bool
    disposing )
{
    if( disposing )
    {
        if (components != null)
        {
            components.Dispose();
        }
        base.Dispose( disposing );
    }
}

#region Windows Form Designer generated code
/// <summary>
/// The main entry point for the application.
/// </summary>
[STAThread]
static void Main()
{
    Application.Run(new Form1());
}

private void button1_Click(object sender,
    System.EventArgs e)
{
    label1.Text = "Hello World";
}
}
}

```

Although this is a simple example it does demonstrate some of the most basic features of C#.NET and should be easily understood by those who have previously developed in C++.

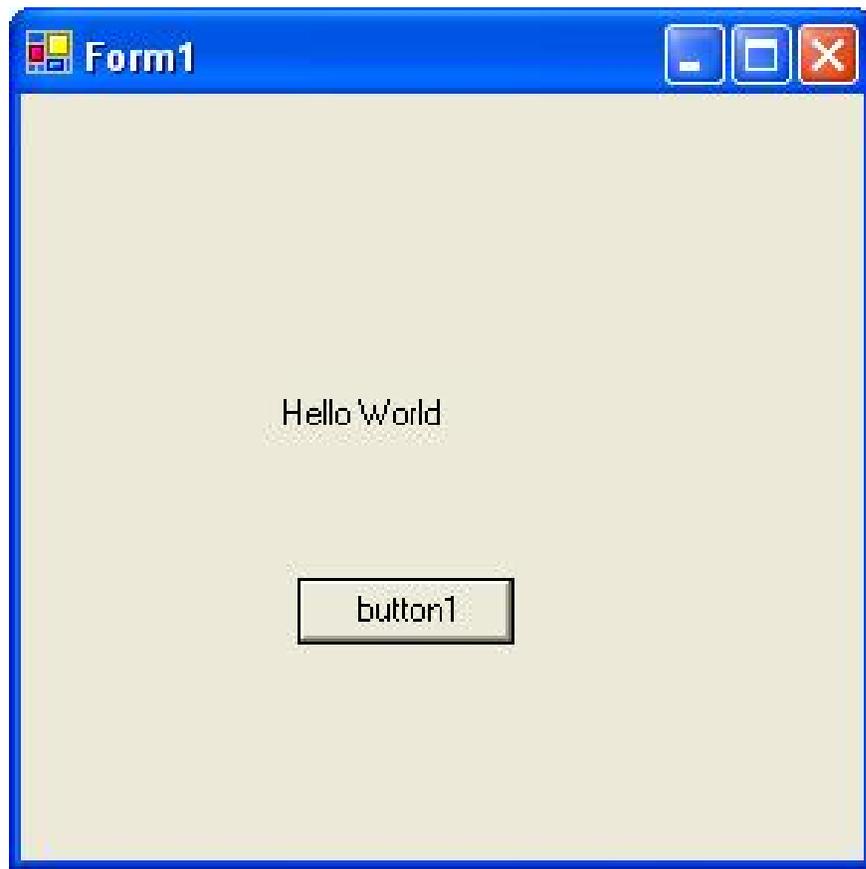


Figure 7.2: The result of Hello World Code execution on the C#.NET form.

7.3 Using Sockets

Understanding the most basic use of C#.NET is only the first step in developing a complex internet application. The next step is to understand the use of Windows Socket. Traditionally a socket was the end-point of a transport layer protocol such as TCP. However a socket became known as the end point for any protocol and the term RAW Socket was coined to describe a socket working below the IP Protocol, that is receiving or transmitting packets with the IP header intact and no IP error checking. (Robison 2002)

7.3.1 Application Programming Interface

In order for a third party application to use a core part of the Operating System it must follow a standard for calling system functions. In windows this standard is called the Application Programming Interface (API). The API controls all sockets as a part of the Operating System function. The socket may be controlled through a number of functions available to the programmer. By calling these functions correctly a program can create a socket connection and send or receive data.

7.3.2 Windows Sockets

The first step in using the API to creating a working socket is to create a socket descriptor. The descriptor is similar to unix files and is usually stored simply as an integer which has special meaning to the operating system. A socket descriptor is created by the code `Socket nameofsocket = null;`. Next, the socket descriptor is linked to a real socket. This is achieved by the code `nameofsocket = new Socket(AddressFamily.InterNetwork, SocketType.Raw, ProtocolType.IP);`. Obviously this socket is for InterNetworks and is a RAW socket from the IP Protocol.

Once a socket is obtained the C#.NET development studio makes it easier by showing the methods implemented for the socket. This interface is shown in Figure 7.3. If the socket command to be accessed is known, then it is only necessary to type the first few

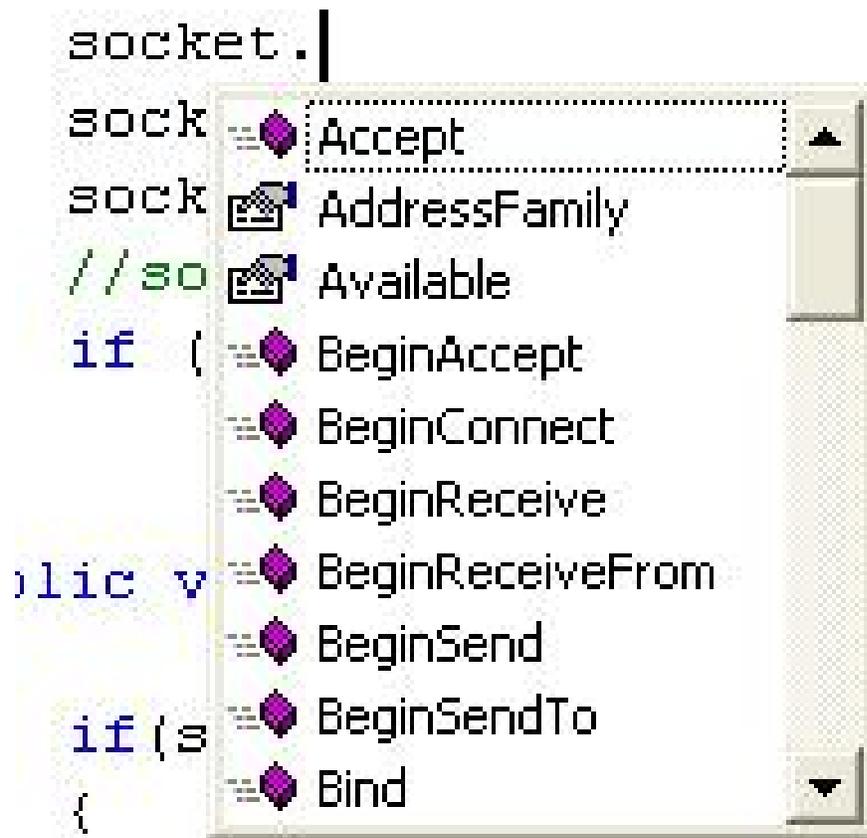


Figure 7.3: The Visual Studio Development Environment.

letters of the name and push the tab key, the Development Studio addresses the rest. It also provides dynamic help based on which command is being used and supports automatic selection for complex or well known choices.

7.3.3 Advanced Socket Control

There are a few special features of sockets required in this project. First there is the requirement to send packets including the IP Header without allowing the system to generate its own IP Header. Secondly, the NAT needs to check all packets as they are received, to ensure that the packets should be sent onto the internet. These packets will not be explicitly sent to the NAT Server because they will contain the address of the remote end-point. This type of receiving is called promiscuous mode and is not supported under all hardware and software configurations.

The solution to the first problem is fairly simple. C#.NET allows an option to be set on each socket to include the IP Header as in the code `nameofsocket.SetSocketOption(SocketOptionLevel.IP, SocketOptionName.HeaderIncluded, 1)`; Setting up Promiscuous mode however, requires a small block of code which requests the appropriate settings and checks the response to ensure the operation completed successfully indicating that the software can support this operation. Listing 7.2 shows the required code segment.

Listing 7.2: Promiscuous Mode Sockets

```
private bool SetSocketOption()
{
    bool ret_value = true;
    try // .NET Exception handling
    {
        byte [] IN = new byte[4]{1, 0, 0, 0};
        byte [] OUT = new byte[4];
        int SIO_RCVALL = unchecked((int)0x98000001);
        // Control code for SIO_RCVALL documented on MSDN.
        // See http://msdn.microsoft.com/library/default.asp?url=/library/en-us/winsock/winsock/wsaiocctl-2.asp for details.
        int ret_code = socket.IOControl(SIO_RCVALL, IN, OUT); // receive all IP packets on the network.
        ret_code = OUT[0] + OUT[1] + OUT[2] + OUT[3];
        // Check that operation succeeded.
        if (ret_code != 0) ret_value = false; // If not return error.
    }
    catch (SocketException)
    {
        ret_value = false; // If any of the above caused an exception, return an error.
    }
    return ret_value;
}
```

Using these advanced features and standard socket operations all the features required for this project were implemented.

7.4 Putting it all Together

Once the Fundamental aspects of working with the RAW Sockets API were understood the final program could be written. This included development of an overview of the solution and finally implementation as a Windows Service.

7.4.1 Pseudocode

Listing 7.3 shows the functioning of the NAT. Obviously this is made very simple in C#.NET by the functionality of Hashing Tables and the RAW Sockets class I wrote which managed all of the IP header and checksum details internally.

Listing 7.3: PseudoCode for Main NAT Function

```
ForEach Packet
    if Packet Source = Internal Network & Packet
        Destination = External Network
        Packet Source = Global Internet Address of NAT
        if Hash Table Result = Port Number
            Packet Source Port = Hash Table Result
        else
            Packet Source Port = New Hash Table Result
    if Packet Destination = Global Internet Address of NAT
        & Destination Port = Reverse Hash Table Result
    Packet Destination = Reverse Hash Table Result
    Packet Destination Port = Original Port Number
```

7.4.2 A Windows Service

Listing 7.4 shows the core code required for an implementation of a Windows Service in C#.NET. It is obvious from this code that the main points are to declare the required variables, set up some constantly looping decision making functions and clean up any variables and persistent code when stopping the service. Attention is drawn to the TODO: labels indicating areas where the user needs to add code. This code is completely generated by C#.NET without any user input except to request the creation of a Windows Service.

Listing 7.4: Implementing a Windows Service

```
| using System;
| using System.Collections;
| using System.ComponentModel;
| using System.Data;
| using System.Diagnostics;
```

```

using System.ServiceProcess;
namespace WindowsService1
{
    public class Service1 : System.ServiceProcess.
        ServiceBase
    {
        /// <summary>
        /// Required designer variable.
        /// </summary>
        private System.ComponentModel.Container
            components = null;

        public Service1()
        {
            // This call is required by the
            // Windows.Forms Component Designer.
            InitializeComponent();

            // TODO: Add any initialization after
            // the InitializeComponent call
        }

        // The main entry point for the process
        static void Main()
        {
            System.ServiceProcess.ServiceBase []
                ServicesToRun;

            // More than one user Service may run
            // within the same process. To add
            // another service to this process,
            // change the following line to
            // create a second service object. For
            // example,
            ///
            /// ServicesToRun = New System.
            /// ServiceProcess.ServiceBase [] {new
            /// Service1(), new
            /// MySecondUserService()};

            ServicesToRun = new System.
                ServiceProcess.ServiceBase [] { new
                Service1() };

            System.ServiceProcess.ServiceBase.Run(
                ServicesToRun);
        }

        /// <summary>
        /// Required method for Designer support - do
        /// not modify
        /// the contents of this method with the code
        /// editor.
        /// </summary>
        private void InitializeComponent()
        {
            components = new System.ComponentModel
                .Container();
            this.ServiceName = "Service1";
        }

        /// <summary>
        /// Clean up any resources being used.
        /// </summary>

```

```
protected override void Dispose( bool
    disposing )
{
    if( disposing )
    {
        if (components != null)
        {
            components.Dispose();
        }
    }
    base.Dispose( disposing );
}

////// <summary>
////// Set things in motion so your service can
////// do its work.
////// </summary>
protected override void OnStart(string[] args)
{
    // TODO: Add code here to start your
    // service.
}

////// <summary>
////// Stop this service.
////// </summary>
protected override void OnStop()
{
    // TODO: Add code here to perform any
    // tear-down necessary to stop your
    // service.
}
}
```

7.5 Chapter Summary

This chapter reviewed some of the significant features of C#.NET that were important in development of the Network Address Translator. By understanding these important concepts the development of the RAW IP Receiver and Sender could be completed. Once this was achieved the main task was to provide an interface for accessing the TCP/IP Header fields for updates and maintaining important information such as the Checksum's updated without involving the user.

The final code implementation is included in Appendix B. There are a few features of this code that have not been discussed here. However the main outcome of this project was a fully features packet class which can be used to receive packets completely, make modifications and send the packet. This class has several other uses in products such as usage meters, network bridges and routers. One important feature of this device is that a router could exist on a network without consuming an IP address while still offering all necessary routing features.

Chapter 8

Conclusions and Further Work

8.1 Achievement of Project Objectives

The following objectives have been addressed:

History of the Internet and Network Reference Models Understanding the motivation behind creating the Internet and the considerations made before its inception, is the first step in developing any type of Internet enabled application. The major focus was on the TCP/IP Protocol suite which is used for most Internet communications. The networking hierarchy was studied in detail to determine the contribution each layer made to the overall communication structure. Two networking hierarchy models were presented, the TCP/IP Reference Model and the Open Systems Interconnect (OSI) Reference Model. These models also discussed common network problems and how they can be overcome. Chapter 2 presented these important details.

Design and System Specification In Chapter 3 the reasons for key design choices were addressed. This project was not intended to be simply thrown away at completion and therefore, needed to adhere to strict communications standards. It was necessary for a design methodology to be chosen to maintain the project time line. Additionally the choice of programming language was discussed and chosen. Although the use of C#.NET was controversial for this type of project

its selection was justified for the special coding features it contained.

The Internet Protocol Chapter 4 expands on a concept raised in the TCP/IP Reference Model from Chapter 2. The idea being to have the layer use a mesh of interconnecting networks to attempt transmission of a packet from a source host through the mesh onto another destination host. At this point Network Address Translator (NAT) implementation begins, however Transport layers also form part of the NAT service delivery and are discussed in Chapter 5.

The Transmission Control Protocol Chapter 5 offers greater detail on the Transport Layer which defines an end-to-end connection between two hosts. The layer provides the error control, retransmission and packet ordering expected by higher level services.

Existing Technology Current implementations of Network Address Translators were discussed in Chapter 6. The importance of utilizing popular features while implementing new ones was also detailed. The costs associated with purchasing some of these existing NATs were also addressed.

Implementation The concepts of C#.NET required in the implementation of a Network Address Translator and other low-level Network applications were addressed in Chapter 7. The topics ranged from a simple greeting program to complex socket operations. Full implementation code is provided in Appendix B.

Although the overall Network Address Translator is incomplete, the major objectives of the project have been achieved. A large amount of research was completed to understand core concepts of Internet communication protocols. Knowledge was gathered concerning features of a chosen programming language appropriate for NAT implementation. The concepts of NAT implementation were comprehensively researched. The code to deal with low level networking interfaces was implemented and the ability to develop NAT in the selected language was proven.

It is anticipated that future project development could use this research and development to implement a fully featured Network Address Translator with security and compatibility features aspired by consumers.

8.2 Further Work

The immediate future of this project requires research into problems experienced with connection stability. The existing code supports full TCP three way handshaking however the connection then resets resulting in client and host confusion. It is suspected that the RAW Socket code may allow the packet to also be passed to higher level layers on the NAT Server which then rejects the connection as unestablished.

In the longer term it would be recommended to implement security and authentication features to ensure only registered users are accessing the NAT server. This could be coupled with packet shaping (slowing the maximum throughput of a client) or connection severing (disconnect clients with no remaining quota). Home users would prefer features that detected expected incoming connections and attempt to compensate at the NAT Server. If such functionality is important sequence numbers could be used to further de-multiplex connections however this could quickly become computationally expensive to the NAT Server and is not a guaranteed method of delivery.

Other features that would integrate well with a NAT project are Firewalls to block unexpected or distrusted inbound connections, Proxy servers to cache web page requests and replies and dynamic port forwarding that open NAT ports when a service such as a Web Server starts and close the ports when the service is stopped.

Testing needs to be undertaken to determine if the NAT and other featured applications would require a multi-threaded approach to ensure maximum user throughput. At the time of development it was considered the development of multiple threads of control within the application were unnecessary. Under heavy load however, a multi-threaded applications may support more clients than a single-threaded application could handle.

References

Advanced Development Methods Inc (2004), *Scrum Development Process*, World Wide Web, United States of America.

<http://www.controlchaos.com/entry.htm>

current May 2004.

Baran, P. (1964), *Distributed Communications*, RAND Corporation.

<http://www.rand.org/publications/RM/RM3420/RM3420.chapter1.html>

current October 2004.

Business ESolutions (2002), *Project Lifecycle Models: How they differ and when to use them*, World Wide Web, California.

<http://www.business-esolutions.com/islm.htm>

current May 2004.

Campione, M., Walrath, K. & Huml, A. (2000), *The Java Tutorial: A Short Course on the Basics*, third edn, Addison-Wesley.

Cerf, V. & Kahn, R. (1974), *A Protocol for Packet Network Interconnection*, Vol. COM-22, IEEE Trans. on Commun., pp. 637–648.

Day, J. & Zimmermann, H. (1983), *The OSI Reference Model*, Vol. 71, The Institute of Electrical and Electronics Engineers, pp. 1334–1340.

Feit, D. S. (1998), *TCP/IP Architecture, Protocols and Implementation with IPv6 and IP Security*, second edn, McGraw-Hill.

Kahn, H. (2000), *STEPWISE Project*, World Wide Web, Manchester.

<http://www.stepwise.org>

current May 2004.

NAT Software (2004), *NAT32 Home Page*, NAT Software Germany.

<http://www.nat32.com/>

current October 2004.

NetcPlus Internet Solutions (2004), *BrowseGate 3 NAT/Proxy server and firewall*,

NetcPlus Internet Solutions.

<http://www.netcplus.com/browsegate.html>

current October 2004.

Perkins, C. E. (2002), *RTP: Audio and Video for the Internet*, Addison-Wesley, Boston.

Rijsinghani, A. (1994), *Computation of the Internet Checksum via Incremental Update*,
Network Working Group.

<ftp://ftp.rfc-editor.org/in-notes/rfc1624.txt>

current October 2004.

Roberts, L. G. (1967), *Multiple Computer Networks and Intercomputer Communication*,
Proc. First Symp. on Operating Systems Prin., ACM.

Robison, W. (2002), *Pure C#: A Code-Intensive Premium Reference*, Sams Publishing.

Sharon, Y. (1999), *(ootips) The Rational Unified Process*, World Wide Web, United
States of America.

<http://ootips.org/rup.html>

current May 2004.

Srisuresh, P. & Egevang, K. (2001), *Traditional IP Network Address Translator (Traditional NAT)*,
Network Working Group.

<ftp://ftp.rfc-editor.org/in-notes/rfc3022.txt>

current October 2004.

Tanenbaum, A. S. (2003), *Computer Networks*, fourth edn, Prentice Hall PTR.

Wells, D. (2003), *Extreme Programming*, World Wide Web, Utah.

<http://www.extremeprogramming.org/index.html>

current May 2004.

Appendix A

Project Specification

University of Southern Queensland

FACULTY OF ENGINEERING AND SURVEYING

ENG 4111/4112 Research Project
PROJECT SPECIFICATION

FOR: **Kevin-John BEASLEY**

TOPIC: Design and Implementation of a Network Address Translator

SUPERVISORS: Dr. John Leis

PROJECT AIM: The project aims to investigate the core aspects of computer networking and TCP/IP Sockets so an efficient and scalable Network Address Translator (NAT) can be designed and implemented for both home and business use.

PROGRAMME: **Issue B, 1 October 2004**

1. Research information on computer networking including history and development.
2. Investigate the TCP/IP Reference Model and TCP/IP Protocols to understand how computer networking relates to the Internet
3. Utilise RFC's relating to NAT Functionality and Implementation.
4. Research appropriate Programming languages for Implementation of a NAT Server, especially the availability of Sockets Programming.
5. Design appropriate sockets interfaces to enable incoming TCP/IP packets to be captured, translated and sent. This includes updating the TCP and IP checksums where necessary.
6. Design a simple NAT server based on the above socket interface to prove the ability to implement NAT in the chosen language.

As time permits:

7. Improve on the NAT implementation to include more features and greater scalability. Support more types of Transport Layer translation.

AGREED: _____ (student) _____ (Supervisor)

____ / ____ / ____ ____ / ____ / ____ (date)

Appendix B

Project Source Code

B.1 NATService.cs

```

1| using System;
2| using System.Collections;
3| using System.ComponentModel;
4| using System.Data;
5| using System.Diagnostics;
6| using System.Net;
7| using System.Net.Sockets;
8| using System.ServiceProcess;
9| using Microsoft.Win32;
10| using System.Security;
11| using System.IO;
12| using System.Extended.Collections;
13|
14| namespace NAT_Service
15| {
16|     public class Service1 : System.ServiceProcess.
17|         ServiceBase
18|     {
19|         /// <summary>
20|         /// Required designer variable and Raw Socket
21|         /// class.
22|         /// </summary>
23|         private System.ComponentModel.Container components
24|             = null;
25|         RawSender myRawSend;
26|         RawSender myIntRawSend;
27|         RawSender[] RawSenders = new RawSender[100];
28|         RawSocket ExternalRawSock;
29|         RawSocket InternalRawSock;
30|         ushort[] ReservedPorts = new ushort[100];
31|         bool[] PortinUse = new bool[100];
32|         BidirHashtable twowayhash = new BidirHashtable();
33|         string ReversedGlobalIP;
34|         string ReversedLocalIP;
35|         int numberofports;
36|
37|     public Service1()
38|     {
39|         // This call is required by the Windows.
40|         // Forms Component Designer.
41|         InitializeComponent();
42|         //TODO: Add any initialization after the
43|         // InitComponent call
44|     }
45|
46|     // The main entry point for the process
47|     static void Main()
48|     {
49|         System.ServiceProcess.ServiceBase[]
50|             ServicesToRun;
51|
52|         // More than one user Service may run within
53|         // the same process. To add
54|         // another service to this process, change
55|         // the following line to
56|         // create a second service object. For
57|         // example,
58|         //
59|         // ServicesToRun = New System.
60|         // ServiceProcess.ServiceBase[] {new
61|         // Service1(), new MySecondUserService()};
62|         //

```

```

53|         ServicesToRun = new System.ServiceProcess.
54|             ServiceBase [] { new Service1 () };
55|         System.ServiceProcess.ServiceBase.Run(
56|             ServicesToRun );
57|     }
58|     /// <summary>
59|     /// Required method for Designer support - do not
60|     /// modify
61|     /// the contents of this method with the code
62|     /// editor.
63|     /// </summary>
64|     private void InitializeComponent ()
65|     {
66|         components = new System.ComponentModel.
67|             Container ();
68|         this.ServiceName = "Service1";
69|     }
70|     /// <summary>
71|     /// Clean up any resources being used.
72|     /// </summary>
73|     protected override void Dispose ( bool disposing )
74|     {
75|         if ( disposing )
76|         {
77|             if ( components != null )
78|             {
79|                 components.Dispose ();
80|             }
81|             base.Dispose ( disposing );
82|         }
83|     }
84|     /// <summary>
85|     /// Set things in motion so your service can do
86|     /// its work.
87|     /// </summary>
88|     protected override void OnStart ( string [] args )
89|     {
90|         // TODO: Add code here to start your service
91|         for ( int i = 0; i < 100; i ++ )
92|         {
93|             ReservedPorts [ i ] = 0;
94|             PortInUse [ i ] = false;
95|         }
96|         RegistryKey key;
97|         try
98|         {
99|             RegistryKey softwareKey = Registry.
100|                 LocalMachine.OpenSubKey ( "Software
101|                 " );
102|             if ( softwareKey == null )
103|             {
104|                 EventLog.WriteEntry ( "Unable to
105|                 open the registry Software
106|                 key for 'USQ_NAT_Project'" );
107|             }
108|             return;
109|         }
110|         key = softwareKey.OpenSubKey ( "
111|             NATUSQProj" );

```

```

105|         if( key == null )
106|         {
107|             EventLog.WriteEntry( "Unable_to_
|                 open_the_registry_NATUSQProj_
|                 _for_'USQ_NAT_Project'" );
108|             return;
109|         }
110|     }
111|     catch( ArgumentNullException argNullExp )
112|     {
113|         EventLog.WriteEntry( "Argument_null_
|                 exception_thrown_" + argNullExp.
|                 Message );
114|         return;
115|     }
116|     catch( ArgumentException argExp )
117|     {
118|         EventLog.WriteEntry( "Argument_
|                 exception_thrown_" + argExp.
|                 Message );
119|         return;
120|     }
121|     catch( IOException ioExp )
122|     {
123|         EventLog.WriteEntry( "IO_Exception_
|                 thrown_" + ioExp.Message );
124|         return;
125|     }
126|     catch( SecurityException secExp )
127|     {
128|         EventLog.WriteEntry( "Security_
|                 exception_thrown_" + secExp.
|                 Message );
129|         return;
130|     }
131|
132|     string LocalIP;
133|     string GlobalIP;
134|     int temp = 0;
135|     ushort testport = 9000;
136|     LocalIP = key.GetValue( "LocalIP" ).ToString
|         ();
137|     GlobalIP = key.GetValue( "GlobalIP" ).
|         ToString();
138|     string [] tempIP;
139|     tempIP = GlobalIP.Split( ".".ToCharArray(), 4)
|         ;
140|     ReversedGlobalIP = tempIP[3] + "." + tempIP
|         [2] + "." + tempIP[1] + "." + tempIP[0];
141|     tempIP = LocalIP.Split( ".".ToCharArray(), 4);
142|     ReversedLocalIP = tempIP[3] + "." + tempIP
|         [2] + "." + tempIP[1] + "." + tempIP[0];
143|     numberofports = (int)key.GetValue( "
|         InitialPorts" );
144|     do
145|     {
146|         bool exception;
147|         do
148|         {
149|             exception = false;
150|             try
151|             {
152|                 myRawSend=new RawSender();

```

```

153|         myRawSend.StartSender (
|             GlobalIP, testport);
154|     }
155|     catch (SocketException ex)
156|     {
157|         if (ex.ErrorCode == 10048)
158|         {
159|             exception = true;
160|         }
161|     }
162|     testport = testport++;
163| } while(exception || myRawSend.WasError
| );
164| RawSenders[temp] = myRawSend;
165| ReservedPorts[temp] = (ushort)(
|     testport - 1);
166| } while(numberofports > temp);
167| myRawSend=new RawSender();
168| myRawSend.StartSender (GlobalIP, 9258);
169| myIntRawSend=new RawSender();
170| myIntRawSend.StartSender (LocalIP, 0);
171| if (myRawSend.WasError)
172| {
173|     EventLog.WriteEntry(" Critical_Error:_
|         Socket_Failed.");
174|     return;
175| }
176| if (myIntRawSend.WasError)
177| {
178|     EventLog.WriteEntry(" Critical_Error:_
|         Socket_Failed.");
179|     return;
180| }
181| string IPString="10.10.10.10";
182| IPHostEntry HosityEntry = Dns.Resolve(Dns.
|     GetHostName());
183| if(HosityEntry.AddressList.Length > 0)
184| {
185|     foreach(IPAddress ip in HosityEntry.
|         AddressList)
186|     {
187|         IPString=ip.ToString();
188|     }
189| }
190|
191| ExternalRawSock=new RawSocket();
192| ExternalRawSock.StartSocket (LocalIP, 0,
|     true);
193| ExternalRawSock.PacketArrival += new
|     RawSocket.PacketArrivedEventHandler(
|     ExternalDataArrival);
194| if (ExternalRawSock.ErrorOccurred)
195| {
196|     EventLog.WriteEntry(" Critical_Error:_
|         Socket_Failed.");
197|     return;
198| }
199| InternalRawSock=new RawSocket();
200| InternalRawSock.StartSocket (GlobalIP, 0,
|     true);
201| InternalRawSock.PacketArrival += new
|     RawSocket.PacketArrivedEventHandler(

```

```

|         InternalDataArrival);
202|     if (InternalRawSock.ErrorOccurred)
203|     {
204|         EventLog.WriteEntry("Critical_Error:_
|             Socket_Failed.");
205|         return;
206|     }
207|     ExternalRawSock.KeepRunning = true; //Want
|         to recieve all incomming packets.
208|     ExternalRawSock.Run ();
209|     InternalRawSock.KeepRunning = true;
210|     InternalRawSock.Run();
211| }
212|
213| private void InternalDataArrival(Object sender,
|     PacketArgs e)
214| {
215|     IPAddress test, test2;
216|     test = new IPAddress((e.source));
217|     test2 = new IPAddress(e.destination);
218|     if((e.destination == (uint)((System.Net.
|         IPAddress.Parse(ReversedGlobalIP).
|         Address))) && (e.tcpdestination == 9258)
|     )
219|     {
220|         ulong returnaddr;
221|         returnaddr = (ulong)twowayhash.
|             ReverseLookup((ulong)((e.
|             destination << 16) + e.
|             tcpdestination));
222|         e.destination = (uint)(returnaddr >>
|             16) & 0xFFFFFFFF;
223|         e.tcpdestination = (ushort)(returnaddr
|             & 0xFFFF);
224|         myIntRawSend.send(e);
225|         if (!myIntRawSend.WasError)
226|         {
227|             // If adding a debugging mode
|                 report that all went well.
228|         }
229|         else
230|         {
231|             EventLog.WriteEntry("Error:_
|                 Return_Packet_Failed.");
232|         }
233|     }
234| }
235|
236| private void ExternalDataArrival(Object sender,
|     PacketArgs e)
237| {
238|     IPAddress test, test2;
239|     test = new IPAddress((e.source));
240|     test2 = new IPAddress(e.destination);
241|     if(e.source == (uint)((System.Net.IPAddress.
|         Parse(ReversedLocalIP).Address)))
242|     {
243|         if (e.tcpdestination == 80)
244|         {
245|             int selectport = -1;
246|             for(int i = 0; i<numberofports -
|                 1; i++)

```



```
285|           // TODO: Add code here to perform any tear-  
|           down necessary to stop your service.  
286|           ExternalRawSock.KeepRunning = false;  
287|           InternalRawSock.KeepRunning = false;  
288|           }  
289|       }  
290| }
```

B.2 RawSocket.cs

```

1| // RawSocket Class
2|
3|
4| namespace NAT_Service
5| {
6|     using System;
7|     using System.Net;
8|     using System.Net.Sockets;
9|     using System.Runtime.InteropServices;
10|
11|
12|     [StructLayout(LayoutKind.Explicit)]
13|     public struct IPHeader
14|     {
15|         [FieldOffset(0)] public byte ip_verIHL; //IP
16|         |           Version (4 bits) + IHL (Header Length) (4 bits
17|         |           )
18|         | //IHL of 5 indicates no options. IHL of 15
19|         | //indicates 40 bytes of options.
20|         [FieldOffset(1)] public byte ip_tos; //Type of
21|         |           Service + Empty (2 bits)
22|         [FieldOffset(2)] public ushort ip_totallength; //
23|         |           Total Packet Length
24|         [FieldOffset(4)] public ushort ip_id; //Unique IP
25|         |           ID
26|         [FieldOffset(6)] public ushort ip_DFMMoffset; //
27|         |           Empty (1 bit) + Don't Fragment (1 bit) + More
28|         |           Fragments (1 bit) Flags + Offset (13 bits)
29|         [FieldOffset(8)] public byte ip_ttl; //Time To
30|         |           Live
31|         [FieldOffset(9)] public byte ip_protocol; //
32|         |           Protocol (TCP, UDP, ICMP, Etc.)
33|         [FieldOffset(10)] public ushort ip_checksum; //IP
34|         |           Header Checksum
35|         [FieldOffset(12)] public uint ip_srcaddr; //
36|         |           Source IP Address
37|         [FieldOffset(16)] public uint ip_destaddr; //
38|         |           Destination IP Address
39|         | // IP Options go here but since we don't know if
40|         | //any exist we won't include them here.
41|     }
42|
43|     [StructLayout(LayoutKind.Explicit)]
44|     public struct TCPHeader
45|     {
46|         [FieldOffset(0)] public ushort tcp_srcport;
47|         |           //Source TCP Port Number.
48|         [FieldOffset(2)] public ushort tcp_destport;
49|         |           //Destination TCP Port Number.
50|         [FieldOffset(4)] public uint tcp_sequence; //TCP
51|         |           Sequence Number
52|         [FieldOffset(8)] public uint tcp_acknowledgement;
53|         |           //TCP Acknowledgement Number;
54|         [FieldOffset(12)] public byte tcp_headerlength; //
55|         |           TCP Header Length (4 bits) + Empty (4 bits)
56|         [FieldOffset(13)] public byte tcp_flags; //Empty
57|         |           (2 bits) + URG (1 bit) + ACK (1 bit) + PSH (1
58|         |           bit) + RST (1 bit) + SYN (1 bit) + FIN (1 bit)
59|         [FieldOffset(14)] public ushort tcp_windowsize
60|         |           ; //TCP Window Size
61|         [FieldOffset(16)] public ushort tcp_checksum; //
62|         |           TCP Checksum

```



```

|         << 8) + buf[IHL * 4 + 11]);
77|     this.TCPHead.tcp_headerlength =
|         buf[IHL * 4 + 12];
78|     this.TCPHead.tcp_flags = buf[IHL
|         * 4 + 13];
79|     this.TCPHead.tcp_windowsize = (
|         ushort)((buf[IHL * 4 + 14]
|         << 8) + buf[IHL * 4 + 15]);
80|     this.TCPHead.tcp_checksum = (
|         ushort)((buf[IHL * 4 + 16]
|         << 8) + buf[IHL * 4 + 17]);
81|     this.TCPHead.tcp_urgentpointer =
|         (ushort)((buf[IHL * 4 + 18]
|         << 8) + buf[IHL * 4 + 19]);
82|     if (this.tcpheaderlength > 5)
83|     {
84|         this.TCPOptions = new byte
|             [(TCPHead.
|             tcp_headerlength - 5)
|             * 4];
85|         Array.Copy(buf, (IHL * 4)
|             + 20, this.TCPOptions,
|             0, (TCPHead.
|             tcp_headerlength - 5)
|             * 4);
86|     }
87|     messagelength = this.MainHeader.
|         ip_totallength - (IHL * 4) -
|         (this.tcpheaderlength * 4);
88|     this.RemainingData = new byte[
|         messagelength];
89|     Array.Copy(buf, (IHL * 4) + (
|         this.tcpheaderlength * 4),
|         this.RemainingData, 0,
|         messagelength);
90|     break;
91|     default: //For any other packets (if
|         implmented) copy everything after
|         the IP Header as data.
92|         messagelength = this.MainHeader.
|             ip_totallength - (IHL * 4);
93|         this.RemainingData = new byte[
|             messagelength];
94|         Array.Copy(buf, IHL * 4, this.
|             RemainingData, 0,
|             messagelength);
95|         break;
96|     }
97| }
98|
99| private UInt16 incrementalchecksum(UInt16 original
|     , UInt16 updated)
100| {
101|     Int32 cksum = Convert.ToInt32(((~(UInt16)
|         this.checksum) & 0xFFFF));
102|     cksum += Convert.ToInt32(((~(UInt16)original
|         ) & 0xFFFF));
103|     cksum += Convert.ToInt32(updated);
104|     cksum = (cksum >> 16) + (cksum & 0xffff);
105|     cksum += (cksum >> 16);
106|     return (UInt16)((~cksum));
107| }

```

```

108|
109|     private UInt16 incrementalTCPchecksum(UInt16
|         original, UInt16 updated)
110|     {
111|         Int32 cksum = Convert.ToInt32(((~(UInt16)
|             this.tcpchecksum) & 0xFFFF));
112|         cksum += Convert.ToInt32(((~(UInt16)original
|             ) & 0xFFFF));
113|         cksum += Convert.ToInt32(updated);
114|         cksum = (cksum >> 16) + (cksum & 0xffff);
115|         cksum += (cksum >> 16);
116|         return (UInt16)(~cksum);
117|     }
118|
119|     /*public static UInt16 calcchecksum( Byte[] buffer
|         , int size )
120|     {
121|         Int32 cksum = 0;
122|         int counter;
123|         counter = 0;
124|         while ( size > 0 )
125|         {
126|             UInt16 val = (ushort)((buffer[counter]
|                 << 8) + buffer[counter+1]);
127|             cksum += Convert.ToInt32( val );
128|             counter += 2;
129|             size -= 2;
130|         }
131|         cksum = (cksum >> 16) + (cksum & 0xffff);
132|         cksum += (cksum >> 16);
133|         return (UInt16)(~cksum);
134|     }*/
135|
136|     public byte version
137|     {
138|         get {return (byte)((this.MainHeader.
|             ip_verIHL & 0xF0) >> 4);}
139|         set
140|         {
141|             ushort temp = (ushort)((this.
|                 MainHeader.ip_verIHL << 8) + this.
|                 MainHeader.ip_tos);
142|             this.MainHeader.ip_verIHL = (byte)((
|                 this.MainHeader.ip_verIHL & 0x0F)
|                 + ((value & 0x0F) << 4));
143|             this.checksum = incrementalchecksum(
|                 temp, (ushort)((this.MainHeader.
|                 ip_verIHL << 8) + this.MainHeader.
|                 ip_tos));
144|         }
145|     }
146|
147|     public byte IHL
148|     {
149|         get {return (byte)(this.MainHeader.ip_verIHL
|             & 0x0F);}
150|         set
151|         {
152|             ushort temp = (ushort)((this.
|                 MainHeader.ip_verIHL << 8) + this.
|                 MainHeader.ip_tos);
153|             ushort tcplength = (ushort)(this.
|                 totallength - (this.IHL * 4));

```

```

154|         this.MainHeader.ip_verIHL = (byte)((
|             this.MainHeader.ip_verIHL & 0xF0)
|             + (value & 0x0F));
155|         this.checksum = incrementalchecksum(
|             temp, (ushort)((this.MainHeader.
|             ip_verIHL << 8) + this.MainHeader.
|             ip_tos));
156|         this.tcpchecksum =
|             incrementalTCPchecksum(tcplength ,
|             (ushort)(this.totallength - (this.
|             IHL * 4)));
157|     }
158| }
159|
160| public byte TOS
161| {
162|     get {return (byte)((this.MainHeader.ip_tos
|         >> 2);}
163|     set
164|     {
165|         ushort temp = (ushort)((this.
|             MainHeader.ip_verIHL << 8) + this.
|             MainHeader.ip_tos);
166|         this.MainHeader.ip_tos = (byte)((value
|             & 0xFF) << 2);
167|         this.checksum = incrementalchecksum(
|             temp, (ushort)((this.MainHeader.
|             ip_verIHL << 8) + this.MainHeader.
|             ip_tos));
168|     }
169| }
170|
171| public ushort totallength
172| {
173|     get {return this.MainHeader.ip_totallength;}
174|     set
175|     {
176|         ushort temp = this.MainHeader.
|             ip_totallength;
177|         ushort tcplength = (ushort)(this.
|             MainHeader.ip_totallength - (this.
|             IHL * 4));
178|         this.MainHeader.ip_totallength = value
|             ;
179|         this.checksum = incrementalchecksum(
|             temp, this.MainHeader.
|             ip_totallength);
180|         this.tcpchecksum =
|             incrementalTCPchecksum(tcplength ,
|             (ushort)(this.MainHeader.
|             ip_totallength - (this.IHL * 4)));
181|     }
182| }
183|
184| public ushort ID
185| {
186|     get {return this.MainHeader.ip_id;}
187|     set
188|     {
189|         ushort temp = this.MainHeader.ip_id;
190|         this.MainHeader.ip_id = value;
191|         this.checksum = incrementalchecksum(
|             temp, this.MainHeader.ip_id);
192|     }

```

```

193|     }
194|
195|     public bool DF
196|     {
197|         get { if((this.MainHeader.ip_DFMMoffset & 0
|             x4000) == 0) return false; else return
|             true; }
198|         set
199|         {
200|             ushort temp = this.MainHeader.
|                 ip_DFMMoffset;
201|             if(value)
202|                 this.MainHeader.ip_DFMMoffset =
|                     (ushort)(this.MainHeader.
|                         ip_DFMMoffset | 0x4000);
203|             else
204|                 this.MainHeader.ip_DFMMoffset =
|                     (ushort)(this.MainHeader.
|                         ip_DFMMoffset & 0x3FFF);
205|             this.checksum = incrementalchecksum(
|                 temp, this.MainHeader.
|                 ip_DFMMoffset);
206|         }
207|     }
208|
209|     public bool MF
210|     {
211|         get { if((this.MainHeader.ip_DFMMoffset & 0
|             x2000) == 0) return false; else return
|             true; }
212|         set
213|         {
214|             ushort temp = this.MainHeader.
|                 ip_DFMMoffset;
215|             if(value)
216|                 this.MainHeader.ip_DFMMoffset =
|                     (ushort)(this.MainHeader.
|                         ip_DFMMoffset | 0x2000);
217|             else
218|                 this.MainHeader.ip_DFMMoffset =
|                     (ushort)(this.MainHeader.
|                         ip_DFMMoffset & 0x5FFF);
219|             this.checksum = incrementalchecksum(
|                 temp, this.MainHeader.
|                 ip_DFMMoffset);
220|         }
221|     }
222|
223|     public ushort Offset
224|     {
225|         get {return (ushort)(this.MainHeader.
|             ip_DFMMoffset & 0x1FF);}
226|         set
227|         {
228|             ushort temp = this.MainHeader.
|                 ip_DFMMoffset;
229|             this.MainHeader.ip_DFMMoffset = (
|                 ushort)((this.MainHeader.
|                     ip_DFMMoffset & 0x7E00) + (value &
|                     0x1FFF));
230|             this.checksum = incrementalchecksum(
|                 temp, this.MainHeader.
|                 ip_DFMMoffset);
231|         }

```

```

232|     }
233|
234|     public byte TTL
235|     {
236|         get {return this.MainHeader.ip_ttl;}
237|         set
238|         {
239|             ushort temp = (ushort)((this.
|                 MainHeader.ip_ttl << 8) + this.
|                 MainHeader.ip_protocol);
240|             this.MainHeader.ip_ttl = value;
241|             this.checksum = incrementalchecksum(
|                 temp, (ushort)((this.MainHeader.
|                 ip_ttl << 8) + this.MainHeader.
|                 ip_protocol));
242|         }
243|     }
244|
245|     public byte Protocol
246|     {
247|         get {return this.MainHeader.ip_protocol;}
248|         set
249|         {
250|             ushort temp = (ushort)((this.
|                 MainHeader.ip_ttl << 8) + this.
|                 MainHeader.ip_protocol);
251|             this.MainHeader.ip_protocol = value;
252|             this.checksum = incrementalchecksum(
|                 temp, (ushort)((this.MainHeader.
|                 ip_ttl << 8) + this.MainHeader.
|                 ip_protocol));
253|             //Could muck around updating the TCP
|                 checksum here but why bother since
|                 protocol != 6 doesn't have a TCP
|                 header anyway.
254|         }
255|     }
256|
257|     private ushort checksum
258|     {
259|         get {return this.MainHeader.ip_checksum;}
260|         set {this.MainHeader.ip_checksum = value;}
261|     }
262|
263|     public uint source
264|     {
265|         get {return this.MainHeader.ip_srcaddr;}
266|         set
267|         {
268|             ushort temp = (ushort)((this.
|                 MainHeader.ip_srcaddr & 0xFFFF0000
|                 ) >> 16);
269|             ushort temp2 = (ushort)(this.
|                 MainHeader.ip_srcaddr & 0x0000FFFF
|                 );
270|             this.MainHeader.ip_srcaddr = value;
271|             this.checksum = incrementalchecksum(
|                 temp, (ushort)((this.MainHeader.
|                 ip_srcaddr & 0xFFFF0000) >> 16));
272|             this.checksum = incrementalchecksum(
|                 temp2, (ushort)(this.MainHeader.
|                 ip_srcaddr & 0x0000FFFF));
273|             this.tcpchecksum =
|                 incrementalTCPchecksum(temp, (

```

```

|         ushort)((this.MainHeader.
|         ip_srcaddr & 0xFFFF0000) >> 16));
274|     this.tcpchecksum =
|         incrementalTCPchecksum(temp2, (
|         ushort)(this.MainHeader.ip_srcaddr
|         & 0x0000FFFF));
|     }
275| }
276|
277|
278| public uint destination
279| {
280|     get {return this.MainHeader.ip_destaddr;}
281|     set
282|     {
283|         ushort temp = (ushort)((this.
|         MainHeader.ip_destaddr & 0
|         xFFFF0000) >> 16);
284|         ushort temp2 = (ushort)(this.
|         MainHeader.ip_destaddr & 0
|         x0000FFFF);
285|         this.MainHeader.ip_destaddr = value;
286|         this.checksum = incrementalchecksum(
|         temp, (ushort)((this.MainHeader.
|         ip_destaddr & 0xFFFF0000) >> 16));
287|         this.checksum = incrementalchecksum(
|         temp2, (ushort)(this.MainHeader.
|         ip_destaddr & 0x0000FFFF));
288|         this.tcpchecksum =
|         incrementalTCPchecksum(temp, (
|         ushort)((this.MainHeader.
|         ip_destaddr & 0xFFFF0000) >> 16));
289|         this.tcpchecksum =
|         incrementalTCPchecksum(temp2, (
|         ushort)(this.MainHeader.
|         ip_destaddr & 0x0000FFFF));
290|     }
291| }
292|
293| public ushort tcpsource
294| {
295|     get {return this.TCPHead.tcp_srcport;}
296|     set
297|     {
298|         ushort temp = this.TCPHead.tcp_srcport
299|         this.TCPHead.tcp_srcport = value;
300|         this.tcpchecksum =
|         incrementalTCPchecksum(temp, this.
|         TCPHead.tcp_srcport);
301|     }
302| }
303|
304| public ushort tcpdestination
305| {
306|     get {return this.TCPHead.tcp_destport;}
307|     set
308|     {
309|         ushort temp = this.TCPHead.
|         tcp_destport;
310|         this.TCPHead.tcp_destport = value;
311|         this.tcpchecksum =
|         incrementalTCPchecksum(temp, this.
|         TCPHead.tcp_destport);
312|     }

```



```

|                                     TCPHead.tcp_flags));
349|     }
350| }
351|
352| public bool tcpUGR
353| {
354|     get {if((this.TCPHead.tcp_flags & 0x20) ==
|         0) return false; else return true; }
355|     set
356|     {
357|         ushort temp = (ushort)((this.TCPHead.
|             tcp_headerlength << 8) + this.
|             TCPHead.tcp_flags);
358|         if(value)
359|             this.TCPHead.tcp_flags = (byte)(
|                 this.TCPHead.tcp_flags | 0
|                 x20);
360|         else
361|             this.TCPHead.tcp_flags = (byte)(
|                 this.TCPHead.tcp_flags & 0
|                 x1F);
362|         this.tcpchecksum =
|             incrementalTCPchecksum(temp, (
|                 ushort)((this.TCPHead.
|                 tcp_headerlength << 8) + this.
|                 TCPHead.tcp_flags));
363|     }
364| }
365|
366| public bool tcpACK
367| {
368|     get {if((this.TCPHead.tcp_flags & 0x10) ==
|         0) return false; else return true; }
369|     set
370|     {
371|         ushort temp = (ushort)((this.TCPHead.
|             tcp_headerlength << 8) + this.
|             TCPHead.tcp_flags);
372|         if(value)
373|             this.TCPHead.tcp_flags = (byte)(
|                 this.TCPHead.tcp_flags | 0
|                 x10);
374|         else
375|             this.TCPHead.tcp_flags = (byte)(
|                 this.TCPHead.tcp_flags & 0
|                 x2F);
376|         this.tcpchecksum =
|             incrementalTCPchecksum(temp, (
|                 ushort)((this.TCPHead.
|                 tcp_headerlength << 8) + this.
|                 TCPHead.tcp_flags));
377|     }
378| }
379|
380| public bool tcpPSH
381| {
382|     get {if((this.TCPHead.tcp_flags & 0x08) ==
|         0) return false; else return true; }
383|     set
384|     {
385|         ushort temp = (ushort)((this.TCPHead.
|             tcp_headerlength << 8) + this.
|             TCPHead.tcp_flags);
```

```

386|         if(value)
387|             this.TCPHead.tcp_flags = (byte)(
|                 this.TCPHead.tcp_flags | 0
|                 x08);
388|         else
389|             this.TCPHead.tcp_flags = (byte)(
|                 this.TCPHead.tcp_flags & 0
|                 x37);
390|         this.tcpchecksum =
|             incrementalTCPchecksum(temp, (
|                 ushort)((this.TCPHead.
|                 tcp_headerlength << 8) + this.
|                 TCPHead.tcp_flags));
391|     }
392| }
393|
394| public bool tcpRST
395| {
396|     get {if((this.TCPHead.tcp_flags & 0x04) ==
|         0) return false; else return true; }
397|     set
398|     {
399|         ushort temp = (ushort)((this.TCPHead.
|             tcp_headerlength << 8) + this.
|             TCPHead.tcp_flags);
400|         if(value)
401|             this.TCPHead.tcp_flags = (byte)(
|                 this.TCPHead.tcp_flags | 0
|                 x04);
402|         else
403|             this.TCPHead.tcp_flags = (byte)(
|                 this.TCPHead.tcp_flags & 0
|                 x3B);
404|         this.tcpchecksum =
|             incrementalTCPchecksum(temp, (
|                 ushort)((this.TCPHead.
|                 tcp_headerlength << 8) + this.
|                 TCPHead.tcp_flags));
405|     }
406| }
407|
408| public bool tcpSYN
409| {
410|     get {if((this.TCPHead.tcp_flags & 0x02) ==
|         0) return false; else return true; }
411|     set
412|     {
413|         ushort temp = (ushort)((this.TCPHead.
|             tcp_headerlength << 8) + this.
|             TCPHead.tcp_flags);
414|         if(value)
415|             this.TCPHead.tcp_flags = (byte)(
|                 this.TCPHead.tcp_flags | 0
|                 x02);
416|         else
417|             this.TCPHead.tcp_flags = (byte)(
|                 this.TCPHead.tcp_flags & 0
|                 x3D);
418|         this.tcpchecksum =
|             incrementalTCPchecksum(temp, (
|                 ushort)((this.TCPHead.
|                 tcp_headerlength << 8) + this.
|                 TCPHead.tcp_flags));

```

```

419|         }
420|     }
421|
422|     public bool tcpFIN
423|     {
424|         get { if((this.TCPHead.tcp_flags & 0x01) ==
425|                |         0) return false; else return true; }
426|         set
427|         {
428|             ushort temp = (ushort)((this.TCPHead.
429|                |         tcp_headerlength << 8) + this.
430|                |         TCPHead.tcp_flags);
431|             if(value)
432|                 this.TCPHead.tcp_flags = (byte)(
433|                    |         this.TCPHead.tcp_flags | 0
434|                    |         x01);
435|             else
436|                 this.TCPHead.tcp_flags = (byte)(
437|                    |         this.TCPHead.tcp_flags & 0
438|                    |         x3E);
439|             this.tcpchecksum =
440|                 incrementalTCPchecksum(temp, (
441|                    |         ushort)((this.TCPHead.
442|                    |         tcp_headerlength << 8) + this.
443|                    |         TCPHead.tcp_flags));
444|         }
445|     }
446|
447|     public ushort tcpwindow
448|     {
449|         get {return this.TCPHead.tcp_windowsize;}
450|         set
451|         {
452|             ushort temp = this.TCPHead.
453|                 tcp_windowsize;
454|             this.TCPHead.tcp_windowsize = value;
455|             this.tcpchecksum =
456|                 incrementalTCPchecksum(temp, this.
457|                 TCPHead.tcp_windowsize);
458|         }
459|     }
460|
461|     private ushort tcpchecksum
462|     {
463|         get {return this.TCPHead.tcp_checksum;}
464|         set {this.TCPHead.tcp_checksum = value;}
465|     }
466|
467|     public ushort tcpurgent
468|     {
469|         get {return this.TCPHead.tcp_urgentpointer;}
470|         set
471|         {
472|             ushort temp = this.TCPHead.
473|                 tcp_urgentpointer;
474|             this.TCPHead.tcp_urgentpointer = value
475|                 ;
476|             this.tcpchecksum =
477|                 incrementalTCPchecksum(temp, this.
478|                 TCPHead.tcp_urgentpointer);
479|         }
480|     }
481| }
482|
483| //private byte proptype;

```

```

465|         public IPHeader MainHeader; //Since I wrote methds
|         for accessing this it should really be
|         private.
466|         // But I don't have time to deal with the changes
|         this requires or to write similar methods for
|         the
467|         //TCP header.
468|         public TCPHeader TCPHead;
469|         public byte [] IPOptions;
470|         public byte [] TCPOptions;
471|         public byte [] RemainingData;
472|         //public int acalculatedchecksum;
473|         public UInt32 test1;
474|         public UInt32 test2;
475|     }
476|
477|     public class RawSender
478|     {
479|         private bool error_in_send;
480|         private static int len_send_buf; //Max packet size
|         for send.
481|         byte [] send_buf_bytes; //Data to send.
482|         private Socket sender = null; //Socket to send via
|
483|         public ushort truecheck;
484|         public ushort reportedcheck;
485|
486|         public RawSender() //Constructor
487|         {
488|             error_in_send=false;
489|             len_send_buf = 65535; //Can't send a larger
|             packet. Windows would crash out.
490|             send_buf_bytes = new byte[len_send_buf];
491|         }
492|
493|         public void StartSender(string IP, int port)
494|         {
495|             sender = new Socket(AddressFamily.
|                 InterNetwork, SocketType.Raw,
|                 ProtocolType.IP);
496|             sender.Blocking = false;
497|             sender.Bind(new IPEndPoint(IPAddress.Parse(
|                 IP), port));
498|             if (SetSenderOption()==false) error_in_send=
|                 true;
499|         }
500|
501|         public void ShutdownSender()
502|         {
503|             if(sender != null)
504|             {
505|                 sender.Shutdown(SocketShutdown.Both);
506|                 sender.Close();
507|             }
508|         }
509|
510|         private bool SetSenderOption()
511|         {
512|             bool ret_value = true;
513|             try //NET Exception handling
514|             {
515|                 sender.SetSocketOption(
|                     SocketOptionLevel.IP,
|                     SocketOptionName.HeaderIncluded,
|                     1);

```

```

516|         sender.SetSocketOption(
|             SocketOptionLevel.Socket, System.
|             Net.Sockets.SocketOptionName.
|             ReuseAddress, 1);
517|         //sender.SetSocketOption(
|             SocketOptionLevel.IP,
|             SocketOptionName.SendBuffer,
|             100000);
518|
519|         //int FIONBIO = unchecked ((int)0
|             x0004CB34);
520|         //int FIONBIO = unchecked ((int)
|             2147772030);
521|         //byte [] IN = new byte[4]{1, 0, 0, 0};
522|         //byte [] OUT = new byte[4];
523|         // See http://msdn.microsoft.com/
|             library/default.asp?url=/library/
|             en-us/winsock/winsock/wsaiocctl-2.
|             asp for details.
524|         //int ret_code = sender.IOControl(
|             FIONBIO, IN, OUT); //receive all
|             IP packets on the network.
525|         //ret_code = OUT[0] + OUT[1] + OUT[2]
|             + OUT[3]; //Check that operation
|             succeeded.
526|         //if (ret_code != 0) ret_value = false;
|             //If not return error.
527|     }
528|     catch(SocketException)
529|     {
530|         ret_value = false; //If any of the
|             above caused an exception, return
|             an error.
531|     }
532|     return ret_value;
533| }
534|
535| public bool WasError
536| {
537|     get // Let main program check for errors.
538|     {
539|         return error_in_send;
540|     }
541| }
542|
543| public static UInt16 calcchecksum( Byte[] buffer ,
|     int size )
544| {
545|     Int32 cksum = 0;
546|     int counter;
547|     counter = 0;
548|     while ( size > 0 )
549|     {
550|         UInt16 val = (ushort)((buffer[counter]
|             << 8) + buffer[counter+1]);
551|         cksum += Convert.ToInt32( val );
552|         counter += 2;
553|         size -= 2;
554|     }
555|     cksum = (cksum >> 16) + (cksum & 0xffff);
556|     cksum += (cksum >> 16);
557|     return (UInt16)(~cksum);
558| }
559|

```

```

560|         public UInt16 TCPChecksum( Byte[] buffer , int size
|         , int IHL )
561|     {
562|         byte[] biggerbuffer;
563|         int length = size - (IHL * 4) + 12;
564|         //this.error_occurred = false;
565|         if (length%2 == 1)
566|         {
567|             //this.error_occurred = true;
568|             length++;
569|         }
570|         biggerbuffer = new byte[length];
571|         //biggerbuffer[0] = buffer[12];
572|         Array.Copy(buffer , 12, biggerbuffer , 0, 8);
573|         biggerbuffer[8] = 0;
574|         biggerbuffer[9] = 6;
575|         biggerbuffer[10] = (byte)((size - (IHL * 4)
|             ) & 0xFF00) >> 8);
576|         biggerbuffer[11] = (byte)((size - (IHL * 4))
|             & 0x00FF);
577|         Array.Copy(buffer , IHL*4, biggerbuffer , 12,
|             size - (IHL * 4));
578|         Int32 cksum = 0;
579|         int counter;
580|         counter = 0;
581|         while ( length > 0 )
582|         {
583|             UInt16 val = (ushort)((biggerbuffer[
|                 counter] << 8) + biggerbuffer[
|                 counter+1]);
584|             cksum += Convert.ToInt32( val );
585|             counter += 2;
586|             length -= 2;
587|         }
588|         cksum = (cksum >> 16) + (cksum & 0xffff);
589|         cksum += (cksum >> 16);
590|         return (UInt16)(~cksum);
591|     }
592|
593|     public void send(PacketArgs args)
594|         //(IPHeader MainHeader)
595|     {
596|         int messagelength;
597|         ushort mychecksum;
598|         //int test = MainHeader.ip_verIHL;
599|         //ListBox1.Items.Add(source + " \t
|             " + destination);
600|         send_buf_bytes[0] = args.MainHeader.
|             ip_verIHL;
601|         send_buf_bytes[1] = args.MainHeader.ip_tos;
602|         send_buf_bytes[2] = (byte)((args.MainHeader.
|             ip_totallength & 0xFF00) >> 8);
603|         send_buf_bytes[3] = (byte)(args.MainHeader.
|             ip_totallength & 0x00FF);
604|         send_buf_bytes[4] = (byte)((args.MainHeader.
|             ip_id & 0xFF00) >> 8);
605|         send_buf_bytes[5] = (byte)(args.MainHeader.
|             ip_id & 0x00FF);
606|         send_buf_bytes[6] = (byte)((args.MainHeader.
|             ip_DFMMoffset & 0xFF00) >> 8);
607|         send_buf_bytes[7] = (byte)(args.MainHeader.
|             ip_DFMMoffset & 0x00FF);

```

```

608|         send_buf_bytes[8] = args.MainHeader.ip_ttl;
609|         send_buf_bytes[9] = args.MainHeader.
|             ip_protocol;
610|         // Don't calculate checksum yet. We want to
|             check both ways until we are certain it
|             works.
611|         send_buf_bytes[10] = 0;
612|         send_buf_bytes[11] = 0;
613|         send_buf_bytes[12] = (byte)((args.MainHeader
|             .ip_srcaddr & 0xFF000000) >> 24);
614|         send_buf_bytes[13] = (byte)((args.MainHeader
|             .ip_srcaddr & 0x00FF0000) >> 16);
615|         send_buf_bytes[14] = (byte)((args.MainHeader
|             .ip_srcaddr & 0x0000FF00) >> 8);
616|         send_buf_bytes[15] = (byte)(args.MainHeader.
|             ip_srcaddr & 0x000000FF);
617|         send_buf_bytes[16] = (byte)((args.MainHeader
|             .ip_destaddr & 0xFF000000) >> 24);
618|         send_buf_bytes[17] = (byte)((args.MainHeader
|             .ip_destaddr & 0x00FF0000) >> 16);
619|         send_buf_bytes[18] = (byte)((args.MainHeader
|             .ip_destaddr & 0x0000FF00) >> 8);
620|         send_buf_bytes[19] = (byte)(args.MainHeader.
|             ip_destaddr & 0x000000FF);
621|         //this.MainHeader.ip_srcaddr = (uint)((buf
|             [12] << 24) + (buf[13] << 16) + (buf[14]
|             << 8) + buf[15]);
622|         int IHL = (args.MainHeader.ip_verIHL & 0x0F)
|             ;
623|         if (IHL > 5)
624|         {
625|             Array.Copy(args.IPOptions, 0,
|                 send_buf_bytes, 20, (IHL - 5) * 4)
|             ;
626|         }
627|         mychecksum = calcchecksum(send_buf_bytes, (
|             send_buf_bytes[0] & 0x0F)*4);
628|         send_buf_bytes[10] = (byte)((args.MainHeader
|             .ip_checksum & 0xFF00) >> 8);
629|         send_buf_bytes[11] = (byte)(args.MainHeader.
|             ip_checksum & 0x00FF);
630|         if (args.Protocol == 6)
631|         {
632|             send_buf_bytes[args.IHL * 4] = (byte)
|                 ((args.TCPHead.tcp_srcport & 0
|                 xFF00) >> 8);
633|             send_buf_bytes[args.IHL * 4 + 1] = (
|                 byte)(args.TCPHead.tcp_srcport & 0
|                 x00FF);
634|             send_buf_bytes[args.IHL * 4 + 2] = (
|                 byte)((args.TCPHead.tcp_destport &
|                 0xFF00) >> 8);
635|             send_buf_bytes[args.IHL * 4 + 3] = (
|                 byte)(args.TCPHead.tcp_destport &
|                 0x00FF);
636|             send_buf_bytes[args.IHL * 4 + 4] = (
|                 byte)((args.TCPHead.tcp_sequence &
|                 0xFF000000) >> 24);
637|             send_buf_bytes[args.IHL * 4 + 5] = (
|                 byte)((args.TCPHead.tcp_sequence &
|                 0x00FF0000) >> 16);

```

```

638|         send_buf_bytes [args.IHL * 4 + 6] = (
|             byte)((args.TCPHead.tcp_sequence &
|                 0x0000FF00) >> 8);
639|         send_buf_bytes [args.IHL * 4 + 7] = (
|             byte)(args.TCPHead.tcp_sequence &
|                 0x000000FF);
640|         send_buf_bytes [args.IHL * 4 + 8] = (
|             byte)((args.TCPHead.
|                 tcp_acknowledgement & 0xFF000000)
|                 >> 24);
641|         send_buf_bytes [args.IHL * 4 + 9] = (
|             byte)((args.TCPHead.
|                 tcp_acknowledgement & 0x00FF0000)
|                 >> 16);
642|         send_buf_bytes [args.IHL * 4 + 10] = (
|             byte)((args.TCPHead.
|                 tcp_acknowledgement & 0x0000FF00)
|                 >> 8);
643|         send_buf_bytes [args.IHL * 4 + 11] = (
|             byte)(args.TCPHead.
|                 tcp_acknowledgement & 0x000000FF);
644|         send_buf_bytes [args.IHL * 4 + 12] =
|             args.TCPHead.tcp_headerlength;
645|         send_buf_bytes [args.IHL * 4 + 13] =
|             args.TCPHead.tcp_flags;
646|         send_buf_bytes [args.IHL * 4 + 14] = (
|             byte)((args.TCPHead.tcp_windowsize
|                 & 0xFF00) >> 8);
647|         send_buf_bytes [args.IHL * 4 + 15] = (
|             byte)(args.TCPHead.tcp_windowsize
|                 & 0x00FF);
648|         send_buf_bytes [args.IHL * 4 + 16] = 0;
649|         send_buf_bytes [args.IHL * 4 + 17] = 0;
650|         send_buf_bytes [args.IHL * 4 + 18] = (
|             byte)((args.TCPHead.
|                 tcp_urgentpointer & 0xFF00) >> 8);
651|         send_buf_bytes [args.IHL * 4 + 19] = (
|             byte)(args.TCPHead.
|                 tcp_urgentpointer & 0x00FF);
652|         if (args.tcpheaderlength > 5)
653|         {
654|             Array.Copy(args.TCPOptions, 0,
|                 send_buf_bytes, (args.IHL *
|                 4) + 20, (args.
|                 tcpheaderlength - 5) * 4);
655|         }
656|         messagelength = args.totallength - (
|             args.IHL * 4) - (args.
|                 tcpheaderlength * 4);
657|         Array.Copy(args.RemainingData, 0,
|             send_buf_bytes, (args.IHL * 4) + (
|                 args.tcpheaderlength * 4),
|             messagelength);
658|     }
659|     else
660|     {
661|         messagelength = args.totallength - (
|             args.IHL * 4);
662|         Array.Copy(args.RemainingData, 0,
|             send_buf_bytes, (args.IHL * 4),
|             messagelength);

```

```

663|     }
664|     //sender.SendTo(send_buf_bytes, args.
|         totallength, System.Net.Sockets.
|         SocketFlags.None, new IPEndPoint(args.
|         destination, 0));
665|     if (mychecksum == args.MainHeader.
|         ip_checksum)
666|     {
667|         if (args.Protocol == 6)
668|         {
669|             mychecksum = TCPChecksum(
|                 send_buf_bytes, args.
|                 totallength, args.IHL);
670|             error_in_send = false;
671|             send_buf_bytes[args.IHL * 4 +
|                 16] = (byte)((args.TCPHead.
|                 tcp_checksum & 0xFF00) >> 8);
672|             send_buf_bytes[args.IHL * 4 +
|                 17] = (byte)(args.TCPHead.
|                 tcp_checksum & 0x00FF);
673|             this.truecheck = mychecksum;
674|             this.reportedcheck = args.
|                 TCPHead.tcp_checksum;
675|             sender.Blocking = false;
676|             //sender.BeginSendTo
677|             sender.Connect(new IPEndPoint((
|                 long)args.destination, args.
|                 tcpdestination));
678|             while (!sender.Connected)
679|             {
680|             }
681|             /*while (sender.Blocking)
682|             {
683|             }*/
684|             sender.Send(send_buf_bytes, 0,
|                 args.totallength, System.Net.
|                 Sockets.SocketFlags.None);
685|             //sender.SendTo(send_buf_bytes,
|                 args.totallength, System.Net.
|                 Sockets.SocketFlags.None,
|                 new IPEndPoint(args.
|                 destination, args.
|                 tcpdestination));
686|         }
687|         else
688|         {
689|             this.truecheck = 0;
690|         }
691|     }
692|     else
693|     {
694|         error_in_send = true;
695|     }
696| }
697|
698| }
699|
700| public class RawSocket
701| {
702|     private bool error_occurred;
703|     public bool KeepRunning; //Keep receiving packets?
704|     private static int len_receive_buf; //Size of
|         receive buffer.

```



```

|
|         details.
748|         int ret_code = socket.IOControl(
|             SIO_RCVALL, IN, OUT); //
|             receive all IP packets on
|             the network.
749|         ret_code = OUT[0] + OUT[1] + OUT
|             [2] + OUT[3]; //Check that
|             operation succeeded.
750|         if (ret_code != 0) ret_value =
|             false; //If not return error
|
|         }
751|     }
752| }
753| catch (SocketException)
754| {
755|     ret_value = false; //If any of the
|         above caused an exception, return
|         an error.
756| }
757| return ret_value;
758| }
759|
760| public bool ErrorOccurred
761| {
762|     get // Let main program check for errors.
763|     {
764|         return error_occurred;
765|     }
766| }
767|
768| public static UInt16 calcchecksum( Byte[] buffer ,
|     int size )
769| {
770|     Int32 cksum = 0;
771|     int counter;
772|     counter = 0;
773|     while ( size > 0 )
774|     {
775|         UInt16 val = (ushort)((buffer[counter]
|             << 8) + buffer[counter+1]);
776|         cksum += Convert.ToInt32( val );
777|         counter += 2;
778|         size -= 2;
779|     }
780|     cksum = (cksum >> 16) + (cksum & 0xffff);
781|     cksum += (cksum >> 16);
782|     return (UInt16)(~cksum);
783| }
784|
785| public UInt16 TCPChecksum( Byte[] buffer , int size
|     , int IHL )
786| {
787|     byte[] biggerbuffer;
788|     int length = size - (IHL * 4) + 12;
789|     //this.error_occurred = false;
790|     if (length%2 == 1)
791|     {
792|         //this.error_occurred = true;
793|         length++;
794|     }
795|     biggerbuffer = new byte[length];
796|     //biggerbuffer[0] = buffer[12];
797|     Array.Copy(buffer , 12, biggerbuffer , 0, 8);
798|     biggerbuffer[8] = 0;

```

```

799|         biggerbuffer[9] = 6;
800|         biggerbuffer[10] = (byte)((((size - (IHL * 4)
|         ) & 0xFF00) >> 8);
801|         biggerbuffer[11] = (byte)((size - (IHL * 4))
|         & 0x00FF);
802|         Array.Copy(buffer, IHL*4, biggerbuffer, 12,
|         size - (IHL * 4));
803|         Int32 cksum = 0;
804|         int counter;
805|         counter = 0;
806|         while ( length > 0 )
807|         {
808|             UInt16 val = (ushort)((biggerbuffer[
|             counter] << 8) + biggerbuffer[
|             counter+1]);
809|             cksum += Convert.ToInt32( val );
810|             counter += 2;
811|             length -= 2;
812|         }
813|         cksum = (cksum >> 16) + (cksum & 0xffff);
814|         cksum += (cksum >> 16);
815|         return (UInt16)(~cksum);
816|     }
817|
818|     private void Receive(byte [] buf, int len)
819|     {
820|         //byte temp_protocol=0;
821|         //uint temp_version=0;
822|         //uint temp_ip_srcaddr=0;
823|         //uint temp_ip_destaddr=0;
824|         //short temp_srcport=0;
825|         //short temp_dstport=0;
826|         //IPAddress temp_ip;
827|
828|         UInt16 RecievedChecksum;
829|         byte temp1, temp2;
830|         RecievedChecksum = (ushort)((buf[10] << 8) +
|         buf[11]);
831|         temp1 = buf[10];
832|         temp2 = buf[11];
833|         buf[10] = 0;
834|         buf[11] = 0;
835|         if (calcchecksum(buf, (buf[0] & 0x0F)*4) ==
|         RecievedChecksum)
836|         {
837|             buf[10] = temp1;
838|             buf[11] = temp2;
839|             int IHL = (buf[0] & 0x0F);
840|             if (buf[9] == 6) //This causes us to
|             only recieve TCP packets which is
|             fine for now.
841|             {
842|                 temp1 = buf[IHL * 4 + 16];
843|                 temp2 = buf[IHL * 4 + 17];
844|                 RecievedChecksum = (ushort)((
|                 temp1 << 8) + temp2);
845|                 buf[IHL * 4 + 16] = 0;
846|                 buf[IHL * 4 + 17] = 0;
847|                 if (TCPChecksum(buf, (ushort)((
|                 buf[2] << 8) + buf[3]), IHL)
|                 == RecievedChecksum)
848|                 {

```

```

849|         buf[IHL * 4 + 16] = temp1;
850|         buf[IHL * 4 + 17] = temp2;
851|         PacketArgs e;
852|         e=new PacketArgs(buf);
853|         OnPacketArrival(e);
854|     }
855|     checksum1=TCPChecksum(buf, (
|         ushort)((buf[2] << 8) + buf
|         [3]), IHL);
856|     checksum2 = RecievedChecksum;
857| }
858| //PacketArgs e;
859|
860| //fixed(byte *fixed_buf = buf)
861| //{
862| //IPHeader * head = (IPHeader *)
|     fixed_buf; // Assign IP Header
|     from recieve buffer.
863| //e=new PacketArgs(buf);
864| // head->ip_protocol, ((uint)(head->
|     ip_verIHL & 0x0F) << 2) * 4
865| //Array.Copy(buf, 0, e.MainHeader, 0,
|     20);
866|
867| //e.MainHeader = head;
868| //e.SubHeader = fixed_buf[((uint)(head
|     ->ip_verIHL & 0x0F) << 2) * 4];
869| /*e.HeaderLength=(uint)(head->
|     ip_verIHL & 0x0F) << 2; //Header
|     Length from IHL (bits 5-8).
870|
871|     temp_protocol = head->
|     ip_protocol;
872|     switch(temp_protocol)
873|     {
874|         case 1: e.Protocol="ICMP
|             ":"; break;
875|         case 2: e.Protocol="IGMP
|             ":"; break; //Don't
|             need IGMP in NAT
876|         case 6: e.Protocol="TCP:";
|             break;
877|         case 17: e.Protocol="UDP
|             ":"; break;
878|         default: e.Protocol= "
|             UNKNOWN"; break;
879|         // See http://www.iana.org
|             /assignments/protocol-
|             numbers for details.
880|     } // Use this in future version
|     to properly decode TCP or
|     ICMP headers.
881|
882|     temp_version =(uint)(head->
|     ip_verIHL & 0xF0) >> 4; //
|     Version from verIHL (bits
|     1-4).
883|     e.IPVersion = temp_version.
|     ToString();
884|
885|     temp_ip_srcaddr = head->
|     ip_srcaddr; //Decode IP
|     addresses to strings.
886|     temp_ip_destaddr = head->
|     ip_destaddr;

```

```

887|         temp_ip = new IPAddress(
|             temp_ip_srcaddr);
888|         e. OriginationAddress =temp_ip.
|             ToString();
889|         temp_ip = new IPAddress(
|             temp_ip_destaddr);
890|         e. DestinationAddress = temp_ip.
|             ToString();
891|
892|         // This is a very bad idea as it
|             defeats the purpose of
|             seperate IP and TCP layers.
893|         // Will reprogram it later when
|             support is added for total
|             recognition of TCP header.
894|         // Would cause problems if IP
|             Options were used. (Could
|             use e.HeaderLength to fix)
895|         temp_srcport = *(short *)&
|             fixed_buf[e.HeaderLength];
896|         temp_dstport = *(short *)&
|             fixed_buf[e.HeaderLength+2];
897|         e. OriginationPort=IPAddress.
|             NetworkToHostOrder(
|             temp_srcport). ToString();
898|         e. DestinationPort=IPAddress.
|             NetworkToHostOrder(
|             temp_dstport). ToString();
899|
900|         e.PacketLength =(uint)len;
901|         e.MessageLength =(uint)len - e.
|             HeaderLength;
902|
903|         e.ReceiveBuffer=buf;
904|         Array.Copy(buf,0,e.
|             IPHeaderBuffer,0,(int)e.
|             HeaderLength);
905|         Array.Copy(buf,(int)e.
|             HeaderLength,e.MessageBuffer
|             ,0,(int)e.MessageLength); //
|             Copy remaining data to
|             message buffers.*/
|
|             //}
906|
907|         //OnPacketArrival(e); //Call
|             processing functions.
908|
909|     }
|
910| }
|
911|
912| public void Run()
913| {
914|     IAsyncResult ar = socket.BeginReceive(
|         receive_buf_bytes , 0, len_receive_buf ,
|         SocketFlags.None, new AsyncCallback(
|         CallReceive) , this);
915| }
916|
917| private void CallReceive(IAsyncResult ar)
918| {
919|     int received_bytes;
920|     received_bytes = socket.EndReceive(ar);
921|     Receive(receive_buf_bytes , received_bytes);
922|     if (KeepRunning) Run(); //Keep recieving
|         more packets.
923| }

```

```
924|
925|
926|     public delegate void PacketArrivedEventHandler(
927|         Object sender, PacketArgs args);
928|
929|     public event PacketArrivedEventHandler
930|         PacketArrival;
931|
932|     protected virtual void OnPacketArrival(PacketArgs
933|         e)
934|     {
935|         if (PacketArrival != null)
936|         {
937|             PacketArrival(this, e);
938|         }
939|     }
```

B.3 BidirHashtable.cs

```

1| using System;
2| using System.Collections;
3| //using System.Runtime.Serialization;
4|
5| namespace System.Extended.Collections
6| {
7|     /// <summary>
8|     /// BidirHashtable is a simple, bidirectional data
9|     /// structure
10|    /// designed around Hashtables and accessed like a more
11|    /// robust Hashtable.
12|    /// Internally it just contains two hashtables:
13|    /// one maps from key to value, the other maps from
14|    /// value to key.
15|    /// Lookup in either direction is quick;
16|    /// changes take twice as long since two Hashtables are
17|    /// accessed.
18|    /// Forward lookup is just through the [] as in
19|    /// Hashtable.
20|    /// Reverse lookup is through ReverseLookup().
21|    /// Adding and setting elements is done with forward
22|    /// syntax identical to
23|    /// in Hashtable, but both internal Hashtables are
24|    /// affected.
25|    /// </summary>
26|    public class BidirHashtable : IDictionary, ICollection,
27|        IEnumerable,
28|        ICloneable
29|    {
30|        private Hashtable m_htFwd = null;
31|        private Hashtable m_htBkwd = null;
32|
33|        /// <summary>
34|        /// Just create a Two-Way Hash Table.
35|        /// </summary>
36|        public BidirHashtable()
37|        {
38|            m_htFwd = new Hashtable();
39|            m_htBkwd = new Hashtable();
40|        }
41|
42|        /// <summary>
43|        /// Somewhat similar to a Copy Constructor in C++
44|        /// </summary>
45|        public BidirHashtable(IDictionary dict)
46|        {
47|            m_htFwd = new Hashtable();
48|            m_htBkwd = new Hashtable();
49|
50|            foreach( object key in dict.Keys )
51|            {
52|                this[key] = dict[key];
53|            }
54|        }
55|
56|        /// <summary>
57|        /// Use an existing Hash Table and map the reverse
58|        /// lookups.
59|        /// </summary>
60|        private BidirHashtable(Hashtable ht, byte
61|            bytDummyIndicatesAttach)
62|        {

```

```

53|         m_htFwd = ht;
54|         m_htBkwd = new Hashtable();
55|
56|         foreach( object key in ht.Keys )
57|         {
58|             m_htBkwd[ht[key]] = key;
59|         }
60|     }
61|
62|     public int Count {get { return m_htFwd.Count; } }
63|     public bool IsSynchronized {get { return m_htFwd.
64|         | IsSynchronized; } }
65|     public object SyncRoot {get { return m_htFwd.
66|         | SyncRoot; } }
67|     public void CopyTo(
68|         Array array,
69|         int index
70|     )
71|     {
72|         m_htFwd.CopyTo( array, index );
73|     }
74|     public void CopyValuesTo(
75|         Array array,
76|         int index
77|     )
78|     {
79|         m_htBkwd.CopyTo( array, index );
80|     }
81|     public void Add( object key, object val )
82|     {
83|         m_htFwd.Add( key, val );
84|         m_htBkwd.Add( val, key );
85|     }
86|
87|     public void Remove( object key )
88|     {
89|         object val = m_htFwd[key];
90|         m_htFwd.Remove( key );
91|         m_htBkwd.Remove( val );
92|     }
93|
94|     public void Clear()
95|     {
96|         m_htFwd.Clear();
97|         m_htBkwd.Clear();
98|     }
99|
100|     /// <summary>
101|     /// Forward lookup or set.
102|     /// </summary>
103|     public object this[ object key ]
104|     {
105|         get { return m_htFwd[key]; }
106|         set
107|         {
108|             if ( m_htFwd.ContainsKey(key) )
109|             {
110|                 m_htBkwd.Remove( m_htFwd[key] );
111|             }
112|             m_htFwd[key] = value;
113|             m_htBkwd[value] = key;
114|         }

```

```

115|     }
116|
117|     /// <summary>
118|     /// Reverse Lookup works at normal hashtable
119|     /// speeds.
120|     /// </summary>
120|     public object ReverseLookup( object val )
121|     {
122|         return m_htBkwd[ val ];
123|     }
124|
125|     public bool IsFixedSize
126|     {
127|         get { return m_htFwd.IsFixedSize; }
128|     }
129|
130|     public bool IsReadOnly
131|     {
132|         get { return m_htFwd.IsReadOnly; }
133|     }
134|
135|     /// <summary>
136|     /// Do Not Use this, should be made private in
137|     /// future.
138|     /// </summary>
138|     public ICollection Keys
139|     {
140|         get { return m_htFwd.Keys; }
141|     }
142|
143|     /// <summary>
144|     /// Do Not Use this, should be made private in
145|     /// future.
146|     /// </summary>
146|     public ICollection Values
147|     {
148|         get { return m_htFwd.Values; }
149|     }
150|
151|     public bool Contains( object key )
152|     {
153|         return m_htFwd.Contains( key );
154|     }
155|
156|     public bool ContainsValue( object val )
157|     {
158|         return m_htBkwd.Contains( val );
159|     }
160|
161|     IEnumerator IEnumerable.GetEnumerator()
162|     {
163|         return m_htFwd.GetEnumerator();
164|     }
165|
166|     IDictionaryEnumerator IDictionary.GetEnumerator()
167|     {
168|         return m_htFwd.GetEnumerator();
169|     }
170|
171|     public object Clone()
172|     {
173|         BidirHashtable bh = new BidirHashtable();
174|         bh.m_htFwd = (Hashtable) m_htFwd.Clone();
175|         bh.m_htBkwd = (Hashtable) m_htBkwd.Clone();
176|         return bh;

```

```

177|     }
178|
179|     #region Explicit conversion to/from Hashtable
180|     public static explicit operator BidirHashtable(
181|         Hashtable ht)
182|     {
183|         return new BidirHashtable( ht );
184|     }
185|     public static explicit operator Hashtable(
186|         BidirHashtable bd)
187|     {
188|         return (Hashtable) bd.m_htFwd.Clone();
189|     }
190| #endregion
191| #region Access to private Hashtables
192| /// <summary>
193| /// Gives direct access for debugging only.
194| /// </summary>
195| public Hashtable ForwardHashtable
196| {
197|     get { return m_htFwd; }
198| }
199|
200| /// <summary>
201| /// Gives direct access for debugging only.
202| /// </summary>
203| public Hashtable BackwardHashtable
204| {
205|     get { return m_htBkwd; }
206| }
207|
208| /// <summary>
209| /// Make a copy to change without causing bugs in
210| /// this
211| /// two-way hash table.
212| /// </summary>
213| public Hashtable BackwardHashtableClone
214| {
215|     get { return (Hashtable) m_htBkwd.Clone(); }
216| }
217| #endregion
218| #region Attach and ReverseDirection
219| /// <summary>
220| /// Reverse the hash table. Lookups are in
221| /// opposite directions.
222| /// </summary>
223| public void ReverseDirection()
224| {
225|     Hashtable htTemp = m_htFwd;
226|     m_htFwd = m_htBkwd;
227|     m_htBkwd = htTemp;
228| }
229| public static BidirHashtable Attach(Hashtable ht)
230| {
231|     return new BidirHashtable( ht, (byte) 0 );
232| }
233| #endregion
234| }
235| }

```

B.4 ProjectInstaller.cs

```

1| using System;
2| using System.Collections;
3| using System.ComponentModel;
4| using System.Configuration.Install;
5| using Microsoft.Win32;
6| using System.Security;
7| using System.IO;
8| using System.Windows.Forms;
9| using System.ServiceProcess;
10| using System.Diagnostics;
11|
12| namespace NAT_Service
13| {
14|     /// <summary>
15|     /// Summary description for ProjectInstaller.
16|     /// </summary>
17|     [RunInstaller(true)]
18|     public class ProjectInstaller : System.Configuration.
19|         Install.Installer
20|     {
21|         private System.ServiceProcess.
22|             ServiceProcessInstaller
23|             serviceProcessInstaller1;
24|         private System.ServiceProcess.ServiceInstaller
25|             serviceInstaller1;
26|
27|         /// <summary>
28|         /// Required designer variable.
29|         /// </summary>
30|         private System.ComponentModel.Container
31|             components = null;
32|         private EventLog eventLog;
33|
34|         public ProjectInstaller()
35|         {
36|             // This call is required by the Designer.
37|             InitializeComponent();
38|             eventLog = new EventLog();
39|
40|             // TODO: Add any initialization after the
41|             // InitComponent call
42|         }
43|
44|         /// <summary>
45|         /// override the install method to set up the
46|         /// information.
47|         /// all thats create here is a registry key. It
48|         /// should be noted that this function
49|         /// can't be debugged so catch all possible
50|         /// exceptions
51|         /// </summary>
52|         /// <param name="iInstallData"></param>
53|         public override void Install( IDictionary
54|             iInstallData )
55|         {
56|             try
57|             {
58|                 // must call base class install first
59|                 base.Install( iInstallData );
60|
61|                 // just create the key the gui part
62|                 // of the code will set it

```

```
53|         RegistryKey reg = Registry.  
|             LocalMachine.OpenSubKey( "Software  
|             ", true );  
54|         if( reg == null )  
55|         {  
56|             eventLog.WriteEntry( "Error_  
|                 trying_to_install_'USQ_NAT_  
|                 Project'" );  
57|             return;  
58|         }  
59|         RegistryKey scheduleKey = reg.  
60|             CreateSubKey( "NATUSQProj" );  
61|         if( scheduleKey == null )  
62|         {  
63|             eventLog.WriteEntry( "Error_  
|                 trying_to_install_'USQ_NAT_  
|                 Project'" );  
64|             return;  
65|         }  
66|         }  
67|         reg.Close();  
68|     }  
69| }  
70| catch( ArgumentNullException argNullExp )  
71| {  
72|     eventLog.WriteEntry( "Error_with_the_  
|         argument_subkey_" + argNullExp.  
|         Message );  
73| }  
74| catch( SecurityException secExp )  
75| {  
76|     eventLog.WriteEntry( "Error_the_user_  
|         does_not_have_access_permission_"  
|         + secExp.Message );  
77| }  
78| catch( IOException ioExp )  
79| {  
80|     eventLog.WriteEntry( "Error_the_  
|         registry_key_is_closed_" + ioExp.  
|         Message );  
81| }  
82| catch( UnauthorizedAccessException unExp )  
83| {  
84|     eventLog.WriteEntry( "Error_the_user_  
|         does_not_have_access_permission_"  
|         + unExp.Message );  
85| }  
86| catch( ArgumentException argExp )  
87| {  
88|     eventLog.WriteEntry( "Error_in_the_  
|         install_data_format_" + argExp.  
|         Message );  
89| }  
90| catch( Exception exp )  
91| {  
92|     eventLog.WriteEntry( "A_problem_  
|         occurred_with_the_install_" + exp.  
|         Message );  
93| }  
94| }  
95| return;  
96| }  
97| }
```

```

98|         /// <summary>
99|         /// override the uninstall method and remove the
100|         /// registry key
101|         /// </summary>
102|         /// <param name="iInstallData"></param>
103|         public override void Uninstall( IDictionary
104|             iInstallData )
105|         {
106|             try
107|             {
108|                 if( iInstallData == null )
109|                 {
110|                     eventLog.WriteEntry( "Error_
111|                         unable_to_uninstall_the_
112|                         application_'USQ_NAT_Project
113|                         '" );
114|                 }
115|                 else
116|                 {
117|                     base.Uninstall( iInstallData );
118|                     Registry.LocalMachine.OpenSubKey
119|                         ( "Software", true ).
120|                         DeleteSubKeyTree( "
121|                             NATUSQProj" );
122|                 }
123|             }
124|             catch( ArgumentException argExp )
125|             {
126|                 MessageBox.Show( "Error_in_the_install
127|                     _data_format_" + argExp.Message );
128|             }
129|             catch( InstallException instExp )
130|             {
131|                 MessageBox.Show( "A_problem_occurred_
132|                     with_the_install_" + instExp.
133|                     Message );
134|             }
135|             return;
136|         }
137|
138| #region Component Designer generated code
139| /// <summary>
140| /// Required method for Designer support - do not
141| /// modify
142| /// the contents of this method with the code
143| /// editor.
144| /// </summary>
145| private void InitializeComponent()
146| {
147|     this.serviceProcessInstaller1 = new System.
148|         ServiceProcess.ServiceProcessInstaller()
149|         ;
150|     this.serviceInstaller1 = new System.
151|         ServiceProcess.ServiceInstaller();
152|
153|     ///
154|     /// serviceProcessInstaller1
155|     ///
156|     this.serviceProcessInstaller1.Account =
157|         System.ServiceProcess.ServiceAccount.
158|         LocalSystem;

```

```
144|         this.serviceProcessInstaller1.Password =
|             null;
145|         this.serviceProcessInstaller1.Username =
|             null;
146|         ///
147|         /// serviceInstaller1
|         ///
148|         this.serviceInstaller1.ServiceName = "
149|             SchedulerService";
150|         this.serviceInstaller1.StartType = System.
|             ServiceProcess.ServiceStartMode.
|             Automatic;
151|         ///
152|         /// ProjectInstaller
|         ///
153|         this.Installers.AddRange(new System.
154|             Configuration.Install.Installer [] {
155|
156|
157|
158|             }
159|         #endregion
160|     }
161| }
```

B.5 NATControl.cs

```

1| using System;
2| using System.Drawing;
3| using System.Collections;
4| using System.ComponentModel;
5| using System.Windows.Forms;
6| using Microsoft.Win32;
7| using System.Data;
8| using System.Net;
9|
10| namespace NAT_Settings_Application
11| {
12|     /// <summary>
13|     /// Summary description for Form1.
14|     /// </summary>
15|     public class NATSetup : System.Windows.Forms.Form
16|     {
17|         private System.Windows.Forms.Label Internal;
18|         private System.Windows.Forms.Label label7;
19|         private System.Windows.Forms.Label label8;
20|         private System.Windows.Forms.Label label9;
21|         private System.Windows.Forms.TextBox
22|             InternalSubnet4;
23|         private System.Windows.Forms.TextBox
24|             InternalSubnet3;
25|         private System.Windows.Forms.TextBox
26|             InternalSubnet2;
27|         private System.Windows.Forms.TextBox
28|             InternalSubnet1;
29|         private System.Windows.Forms.Label label10;
30|         private System.Windows.Forms.Label label11;
31|         private System.Windows.Forms.GroupBox groupBox1;
32|         private System.Windows.Forms.GroupBox groupBox2;
33|         private System.Windows.Forms.ComboBox ExternalIP;
34|         private System.Windows.Forms.ComboBox InternalIP;
35|         private System.Windows.Forms.Button NATOkButton;
36|         private System.Windows.Forms.Button
37|             NATCancelButton;
38|         private System.Windows.Forms.Label label1;
39|         private System.Windows.Forms.TextBox NumberOfPorts;
40|
41|         /// <summary>
42|         /// Required designer variable.
43|         /// </summary>
44|         private System.ComponentModel.Container components
45|             = null;
46|
47|         public NATSetup()
48|         {
49|             //
50|             // Required for Windows Form Designer
51|             // support
52|             //
53|             InitializeComponent();
54|
55|             //
56|             // TODO: Add any constructor code after
57|             // InitializeComponent call
58|             //
59|         }
60|
61|         /// <summary>
62|         /// Clean up any resources being used.
63|         /// </summary>
64|         protected override void Dispose( bool disposing )

```

```
56|         {
57|             if( disposing )
58|             {
59|                 if (components != null)
60|                 {
61|                     components.Dispose();
62|                 }
63|             }
64|             base.Dispose( disposing );
65|         }
66|
67| #region Windows Form Designer generated code
68| /// <summary>
69| /// Required method for Designer support - do not
70| /// modify
71| /// the contents of this method with the code
72| /// editor.
73| /// </summary>
74| private void InitializeComponent()
75| {
76|     this.Internal = new System.Windows.Forms.
77|         Label();
78|     this.label7 = new System.Windows.Forms.Label
79|         ();
80|     this.label8 = new System.Windows.Forms.Label
81|         ();
82|     this.label9 = new System.Windows.Forms.Label
83|         ();
84|     this.InternalSubnet4 = new System.Windows.
85|         Forms.TextBox();
86|     this.InternalSubnet3 = new System.Windows.
87|         Forms.TextBox();
88|     this.InternalSubnet2 = new System.Windows.
89|         Forms.TextBox();
90|     this.InternalSubnet1 = new System.Windows.
91|         Forms.TextBox();
92|     this.label10 = new System.Windows.Forms.
93|         Label();
94|     this.label11 = new System.Windows.Forms.
95|         Label();
96|     this.groupBox1 = new System.Windows.Forms.
97|         GroupBox();
98|     this.InternalIP = new System.Windows.Forms.
99|         ComboBox();
100|     this.groupBox2 = new System.Windows.Forms.
101|         GroupBox();
102|     this.ExternalIP = new System.Windows.Forms.
103|         ComboBox();
104|     this.NATOkButton = new System.Windows.Forms.
105|         Button();
106|     this.NATCancelButton = new System.Windows.
107|         Forms.Button();
108|     this.label1 = new System.Windows.Forms.Label
109|         ();
110|     this.NumberofPorts = new System.Windows.
111|         Forms.TextBox();
112|     this.groupBox1.SuspendLayout();
113|     this.groupBox2.SuspendLayout();
114|     this.SuspendLayout();
115|     ///
116|     /// Internal
117|     ///
```

```
98|         this.Internal.Location = new System.Drawing.  
|             Point(32, 32);  
99|         this.Internal.Name = "Internal";  
100|         this.Internal.Size = new System.Drawing.Size  
|             (64, 16);  
101|         this.Internal.TabIndex = 14;  
102|         this.Internal.Text = "IP_Address";  
103|         this.Internal.TextAlign = System.Drawing.  
|             ContentAlignment.MiddleRight;  
104|         ///  
105|         /// label7  
106|         ///  
107|         this.label7.Location = new System.Drawing.  
|             Point(192, 72);  
108|         this.label7.Name = "label7";  
109|         this.label7.Size = new System.Drawing.Size  
|             (8, 16);  
110|         this.label7.TabIndex = 21;  
111|         this.label7.Text = ".";  
112|         ///  
113|         /// label8  
114|         ///  
115|         this.label8.Location = new System.Drawing.  
|             Point(160, 72);  
116|         this.label8.Name = "label8";  
117|         this.label8.Size = new System.Drawing.Size  
|             (8, 16);  
118|         this.label8.TabIndex = 20;  
119|         this.label8.Text = ".";  
120|         ///  
121|         /// label9  
122|         ///  
123|         this.label9.Location = new System.Drawing.  
|             Point(128, 72);  
124|         this.label9.Name = "label9";  
125|         this.label9.Size = new System.Drawing.Size  
|             (8, 16);  
126|         this.label9.TabIndex = 19;  
127|         this.label9.Text = ".";  
128|         ///  
129|         /// InternalSubnet4  
130|         ///  
131|         this.InternalSubnet4.Location = new System.  
|             Drawing.Point(200, 72);  
132|         this.InternalSubnet4.MaxLength = 3;  
133|         this.InternalSubnet4.Name = "InternalSubnet4  
|             ";  
134|         this.InternalSubnet4.Size = new System.  
|             Drawing.Size(24, 20);  
135|         this.InternalSubnet4.TabIndex = 18;  
136|         this.InternalSubnet4.Text = "";  
137|         ///  
138|         /// InternalSubnet3  
139|         ///  
140|         this.InternalSubnet3.Location = new System.  
|             Drawing.Point(168, 72);  
141|         this.InternalSubnet3.MaxLength = 3;  
142|         this.InternalSubnet3.Name = "InternalSubnet3  
|             ";  
143|         this.InternalSubnet3.Size = new System.  
|             Drawing.Size(24, 20);  
144|         this.InternalSubnet3.TabIndex = 17;  
145|         this.InternalSubnet3.Text = "";  
146|         ///  
147|         /// InternalSubnet2
```

```
148| //
149| this.InternalSubnet2.Location = new System.
|     Drawing.Point(136, 72);
150| this.InternalSubnet2.MaxLength = 3;
151| this.InternalSubnet2.Name = "InternalSubnet2
|     ";
152| this.InternalSubnet2.Size = new System.
|     Drawing.Size(24, 20);
153| this.InternalSubnet2.TabIndex = 16;
154| this.InternalSubnet2.Text = "";
155| //
156| // InternalSubnet1
157| //
158| this.InternalSubnet1.Location = new System.
|     Drawing.Point(104, 72);
159| this.InternalSubnet1.MaxLength = 3;
160| this.InternalSubnet1.Name = "InternalSubnet1
|     ";
161| this.InternalSubnet1.Size = new System.
|     Drawing.Size(24, 20);
162| this.InternalSubnet1.TabIndex = 15;
163| this.InternalSubnet1.Text = "";
164| //
165| // label10
166| //
167| this.label10.Location = new System.Drawing.
|     Point(24, 72);
168| this.label10.Name = "label10";
169| this.label10.Size = new System.Drawing.Size
|     (72, 16);
170| this.label10.TabIndex = 22;
171| this.label10.Text = "Subnet_Mask";
172| this.label10.TextAlign = System.Drawing.
|     ContentAlignment.MiddleRight;
173| //
174| // label11
175| //
176| this.label11.Location = new System.Drawing.
|     Point(32, 32);
177| this.label11.Name = "label11";
178| this.label11.Size = new System.Drawing.Size
|     (64, 16);
179| this.label11.TabIndex = 23;
180| this.label11.Text = "IP_Address";
181| this.label11.TextAlign = System.Drawing.
|     ContentAlignment.MiddleRight;
182| //
183| // groupBox1
184| //
185| this.groupBox1.Controls.AddRange(new System.
|     Windows.Forms.Control[] {
186|         this.InternalIP ,
187|         this.label7 ,
188|         this.label8 ,
189|         this.label9 ,
190|         this.InternalSubnet3 ,
191|         this.InternalSubnet2 ,
192|         this.InternalSubnet4 ,
193|         this.label10 ,
194|         this.InternalSubnet1 ,
195|         this.Internal});
196| this.groupBox1.Location = new System.Drawing
|     .Point(56, 24);
197| this.groupBox1.Name = "groupBox1";
```



```

247|         this.label1.Location = new System.Drawing.
|             Point(64, 280);
248|         this.label1.Name = "label1";
249|         this.label1.Size = new System.Drawing.Size
|             (128, 24);
250|         this.label1.TabIndex = 39;
251|         this.label1.Text = "Number_of_Ports_to_Use:_
|             _";
252|         this.label1.TextAlign = System.Drawing.
|             ContentAlignment.MiddleRight;
253|         ///
254|         /// NumberofPorts
255|         ///
256|         this.NumberofPorts.Location = new System.
|             Drawing.Point(200, 280);
257|         this.NumberofPorts.Name = "NumberofPorts";
258|         this.NumberofPorts.Size = new System.Drawing
|             .Size(112, 20);
259|         this.NumberofPorts.TabIndex = 38;
260|         this.NumberofPorts.Text = "10";
261|         ///
262|         /// NATSetup
263|         ///
264|         this.AutoScaleBaseSize = new System.Drawing.
|             Size(5, 13);
265|         this.ClientSize = new System.Drawing.Size
|             (376, 382);
266|         this.Controls.AddRange(new System.Windows.
|             Forms.Control[] {
267|             this.label1,
268|             this.NumberofPorts,
269|             this.NATCancelButton,
270|             this.NATOkButton,
271|             this.groupBox2,
272|             this.groupBox1});
273|         this.Name = "NATSetup";
274|         this.Text = "NAT_Setup_and_Control";
275|         this.Load += new System.EventHandler(this.
|             NATSetup_Load);
276|         this.groupBox1.ResumeLayout(false);
277|         this.groupBox2.ResumeLayout(false);
278|         this.ResumeLayout(false);
279|     }
280| }
281| #endregion
282|
283| /// <summary>
284| /// The main entry point for the application.
285| /// </summary>
286| [STAThread]
287| static void Main()
288| {
289|     Application.Run(new NATSetup());
290| }
291|
292| private void NATSetup_Load(object sender, System.
|     EventArgs e)
293| {
294|     IPEndPoint HosityEntry = Dns.Resolve(Dns.
|         GetHostName());
295|     if (HosityEntry.AddressList.Length > 0)
296|     {
297|         foreach (IPAddress ip in HosityEntry.
|             AddressList)

```

```

298|         {
299|             ExternalIP.Items.Add(ip.ToString
|             ());
300|             InternalIP.Items.Add(ip.ToString
|             ());
301|         }
302|     }
303| }
304|
305| private void NATCancelButton_Click(object sender,
|     System.EventArgs e)
306| {
307|     Application.Exit();
308| }
309|
310| private void NATOkButton_Click(object sender,
|     System.EventArgs e)
311| {
312|     /// safety checking code
313|     if( InternalIP.SelectedIndex == -1 &&
|         ExternalIP.SelectedIndex == -1 )
314|     {
315|         MessageBox.Show( "You_must_select_an_
|             Internal_and_External_IP_address"
|             );
316|         return;
317|     }
318|     if (!( (System.Convert.ToInt32(NumberOfPorts.
|         Text) >= 1) && (System.Convert.ToInt32(
|         NumberOfPorts.Text) <= 60000)))
319|     {
320|         MessageBox.Show( "You_must_enter_a_
|             number_of_ports_to_use..Currently_
|             Min=-1_and_Max=-60000" );
321|         return;
322|     }
323|     /// Edit the Registry.
324|     RegistryKey reg = Registry.LocalMachine.
|         OpenSubKey( "Software", true ).
|         OpenSubKey( "NATUSQProj", true );
325|     if( reg == null )
326|     {
327|         MessageBox.Show( "Error_unable_to_
|             create_the_registry_key_'USQ_NAT_
|             Project'" );
328|         return;
329|     }
330|     reg.SetValue( "LocalIP", InternalIP.Text);
331|     reg.SetValue( "GlobalIP", ExternalIP.Text);
332|     reg.SetValue( "InitialPorts", (NumberOfPorts
|         .Text));
333|
334|     Application.Exit();
335| }
336| }
337| }

```