Computer Science and Artificial Intelligence Laboratory
Technical Report

MIT-CSAIL-TR-2006-004 January 11, 2006

Service Identification in TCP/IP: Well-Known versus Random Port Numbers

Elizabeth Masiello

massachusetts institute of technology, cambridge, ma 02139 usa — www.csail.mit.edu

Service Identification in TCP/IP: Well-Known versus Random Port Numbers

by

Elizabeth Masiello

B.A., Computer Science
Wellesley College 2003

Submitted to the Engineering Systems Division
in Partial Fulfillment of the Requirements for the Degree of

Master of Science in Technology and Policy

at the

Massachusetts Institute of Technology

September 2005

Signature of Author……………………………………………………………………………..
Technology and Policy Program, Engineering Systems Division
May 24, 2005

Certified by……………………………………………………………………………………..
David D. Clark
Senior Research Scientist
Thesis Supervisor

Accepted by……………………………………………………………………………………..
Dava J. Newman
Professor of Aeronautics and Astronautics and Engineering Systems
Director, Technology and Policy Program

Service Identification in TCP/IP:
Well-Known versus Random Port Numbers

by

Elizabeth Masiello

Submitted to the Engineering Systems Division on May 24, 2005 in Partial Fulfillment of the Requirements for
the Degree of Master of Science in Technology and Policy

Abstract

The sixteen-bit well-known port number is often overlooked as a network identifier in Internet communications. Its purpose at the most fundamental level is only to demultiplex flows of traffic. Several unintended uses of the port number evolved from associating services with a list of well-known port numbers. This thesis documents those unintended consequences in an effort to describe the port number's influence on Internet players from ISPs to application developers and individual users. Proposals and examples of moving away from well-known port numbers to randomly assigned ones are then presented, with analysis of impacts on the political and economic systems on which Internet communication is dependent.

Thesis Supervisor: David D. Clark
Senior Research Scientist

Acknowledgements

My most sincere gratitude is extended to Dr. David Clark. Without his patience and guidance, this work would not have been possible, nor would my time at MIT have been nearly as manageable. He has not only been an incredible advisor but also a remarkable teacher.

I was fortunate to have more than one teacher through this process. Extensive gratitude is owed to Steve Bauer and Peyman Faratin who not only withstood having me as an officemate, but also patiently explained networking concepts and details whenever asked.

Finally, many thanks are owed to the friends and family that supported me during these two years. After all the work is completed, I will most cherish our fun and shared memories at MIT.

Table of Contents

Chapter 1 Introduction

A key challenge facing Internet architects is anticipating the unintended consequences of their design decisions. In network design, seemingly minute details have had vast consequences as the scale and ubiquity of the Internet have grown over time. Perhaps the best-known example of this is the decision to make Internet addresses 32 bits, capping the number of addressable machines at a number that seemed unnecessarily large at the time but has since been perceived small enough to warrant development of an entirely new address space.[1] Had the initial design incorporated the uncertainty around network expansion, transition to a larger address space might have been less complicated or difficult to achieve.

The purpose of this thesis is to examine the unintended consequences that resulted from a similarly innocuous design decision, using well-known port numbers, and attempt to anticipate the consequences of an alternative design. There are two primary objectives. The first of these is to document the unintended consequences of using well-known port numbers to identify applications and protocols. The second is to provide a rigorous analysis of the interests in hiding or revealing identity information as represented by the port number, with specific attention to how political, economic and technical power might be shifted by implementing random port numbers in future architectures.

I. The Question

The port number was originally included in Internet architecture as a means of demultiplexing chunks of data into coherent streams of data in a packet-switched network. In this respect the port number corresponds to a process at an endpoint and is used to direct packets to the appropriate system processes. The port number acquired a second meaning as a universal application identifier, the result of a seemingly innocuous decision that has subsequently had an array of consequences across the network.

To establish communication with a service, a client has to know the corresponding port number to which requests should be sent. The implementation chosen to facilitate this, most likely the simplest implementation was a list of well-known port numbers. As a result the port number has acquired additional meaning as an identifier of application or protocol.

The simple decision to use well-known port numbers has had a set of unintended consequences throughout the network. The port number has become central to several methods used by Internet Service Providers (ISPs) in traffic analysis. ISPs also use the port number to create pricing stratification by disabling or enabling specific services for different customer segments. Network security has in some of its most basic elements a reliance on the port number: the most common firewalls are built on the premise that unauthorized traffic can be blocked based on the service with which it is associated.

Port numbers could serve the purpose of demultiplexing without corresponding to a predetermined and well-known list of protocols and applications. As early as the 1980s design alternatives using "rendezvous strings" and random port numbers were articulated and juxtaposed with the decision to use well-known port numbers:

> "…the "well-known port" problem can be solved by devising a second mechanism, distinct from port dispatching, to name well-known ports. Several protocols have settled on the idea of including…a more general service descriptor, such as a character string field. These special packets, which are requesting connection to a particular service, are routed on arrival to a special server, sometimes called a 'rendezvous server', which…selects a random port which is to be used for this instance of the service…"[16]

[1] IPV6 was developed primarily to increase the space of available IP addresses.

This early description of using random port numbers foresaw developing preferences for alternative designs. In recent years new application protocols have begun using random port numbers, indicating preference for hiding the application identity from third-party observers. As next-generation Internet architectures are designed, the conflicting preferences for revealed or hidden identity warrant evaluation.

II. The Methodology

This thesis frames the question of random versus well-known ports in the context of tussle as articulated by Clark *et al.*[17] Engineers influence systems through architecture design. Effective influence is achieved by understanding conflicting interests and designing the space in which competition can take place. Understanding the interests and actions taken to influence use of well-known port numbers as identifiers will clarify obstacles and opportunities facing designers building future network architectures.

This thesis is an early attempt at rigorous analysis of the tussle space in which application identity sits on the network. A rigid methodology for such an analysis does not yet exist but is close in nature to policy analysis. Because it is forward-looking and relies on abstract value definitions, the analysis cannot rely on data currently available, as many social science analyses might.

The methodology whose objectives most closely match those here is a stakeholder analysis. Stakeholder analyses are often aimed at understanding the interests of each party involved and the power held by each to achieve their objectives. Here, the interests of each party are the key focus with secondary attention paid to power. The primary goal is to inform architecture design by identifying, documenting and projecting current and potential interests in use of the port number.

Chapter 2 Related Work

This thesis represents the first documentation of political and economic interests relying on the sixteen-bit port number. As such, it builds on work from a variety of disciplines and on several technical Internet features, none of which is wholly related.

The foundation of this thesis is work in the network architecture space. This includes the articulation by Saltzer, et. al. [66] on design principals that shaped today's Internet architecture, as well as recommendations for designing future architectures, most notably Clark's articulation of the Tussle.[17] The direct foundation of this thesis is the early articulation by Clark that well-known ports were only one of several alternative designs, and the uncertainty around whether it was in fact the most appropriate.[16]

Related to network architecture is the space of analysis on adapting the Internet to current political and economic environments. This includes but is not limited to work by Blumenthal et. al.[6] and Lemley, et. al.[47] on the future of end-to-end arguments. Also related, though more remotely, are principals of network neutrality as discussed by Wu[79] and Yoo[81].

The technical underpinnings of the Internet architecture are directly relevant, including but not limited to IP, TCP, DNS and NAT specifications and implementations.[60][61][53][54][25] Transport protocols other than TCP exist, though are not as widely used today and are not the focus here. These include UDP, CHAOSnet, Xerox NS and DECnet.[62][55][80][77]CHAOSnet and XNS used ports differently than TCP/IP do today; CHAOSnet implemented random port numbers using a "rendezvous string" while XNS moved the port number into the network layer rather than the transport layer.

Also centrally important to this thesis is work on implementing random ports. Some of this work has come from the research community, most notably the DNS SRV records.[36] Several applications are implementing random ports in real-time, and though the details of implementation are often proprietary, documentation of their existence is relevant. Clark et.al.[18] articulate design specifications for an architecture using random ports and also implement a prototype.

Work in the area of traffic analysis and packet identification is also related, for example by Fraleigh et. al.[32] and Karagiannis[42]. Cheswick et.al. discuss packet identification in the context of network security and firewalls, which interact with the port number.[13]

This thesis is the first attempt at documenting how the port number interacts with and influences all aspects of the Internet. While related work forms the foundation of this document, there are no known papers dealing with just the port number. The closest such work is a section of Solum et.al.'s work on network layers and the law that discusses port blocking and P2P, though this discussion is limited to the P2P case.[71] Further work in this area would build a more rigorous analysis comparing the TCP/IP protocol's use of port numbers to that of other, less widely deployed protocols.

Chapter 3 Background

This chapter introduces Internet architecture with specific attention to the sixteen-bit port number. It will provide a brief technical overview somewhat limited in scope because of its focus on the port number. The technical details explained in this chapter lay the groundwork for the core analysis of this work.

I. Internet Basics

Before focusing on how and why port numbers are used in Internet communication, it is necessary to review a few key concepts and terms. The Internet is in its most basic definition a network of computers, each of which is an *endpoint* of the network. In between the endpoints are *routers* that route information through the network from source to destination. The original architecture design pushed functionality to the endpoints of the network, making the switches the "dumb" internals responsible only for routing.[2]

The Internet is a packet-switched network, meaning that switches forward discrete blocks of data in the form of *packets*. Each packet has a header and a body; the header tells routers where to send the packet, while the body contains the data content. When data is sent across the network, it is broken into a sequence of packets which must be reassembled at the destination. A coherent sequence of packets makes up a *flow* of traffic. Packets are *multiplexed*, meaning that multiple packets, and therefore multiple flows, can travel simultaneously across a given link on the network. An endpoint that receives multiple packets at once must *demulitplex* those packets into their respective flows.

The Internet has evolved such that most communications are *client-server*, though a growing amount of traffic is *peer-to-peer* (P2P). In the client-server model, a client (or host) will send a request to a server. The server listens for incoming requests and responds appropriately. As an example, Web servers such as the endpoint running http://web.mit.edu listen for incoming requests from clients such as your home computer. In a P2P model, both endpoints are capable of listening for or sending requests. In other words, endpoints in the P2P model act as both clients and servers.

Each endpoint is addressed using the Internet Protocol (IP) which assigns endpoints IP numbers. IP numbers can be looked up globally through a Domain Name System (DNS) that provides address look-up tables. The IP number was originally designed to be 32 bits long allowing for a finite four billion possible IP numbers, a number originally perceived to be more than large enough to accommodate the needs of the fledgling Internet.

As the Internet evolved and grew to achieve global scale and reach, the finite address space became perceived as scarce as opposed to vast. Primarily in response to a perceived shortage, engineers developed Network Address Translation (NAT) that has since become an integral component of the Internet.[25] NAT enables translation from one publicly addressable IP address to a set of local, privately addressable IP addresses, extending the available IP address space. Though some argue that NAT breaks the end-to-end design by inserting functionality (in the form of address translation) into the inside of the network, NAT has been accepted as a permanent component of the network.

[2] This design was articulated most clearly by Saltzer, Reed and Clark and is commonly referred to as the end-to-end design.[66]

II. TCP/IP and Port Numbers

Most Internet communication is done using the Transmission Control Protocol and Internet Protocol (TCP/IP), a common set of protocols for routing packets. TCP is the protocol used at the endpoints to disassemble data into packets for delivery and reassemble packets into data for reception. IP is used by routers to forward packets from origin to destination. TCP and IP are implemented by attaching header information onto packets sent across the network. The TCP and IP headers are displayed in Figures 3-1.

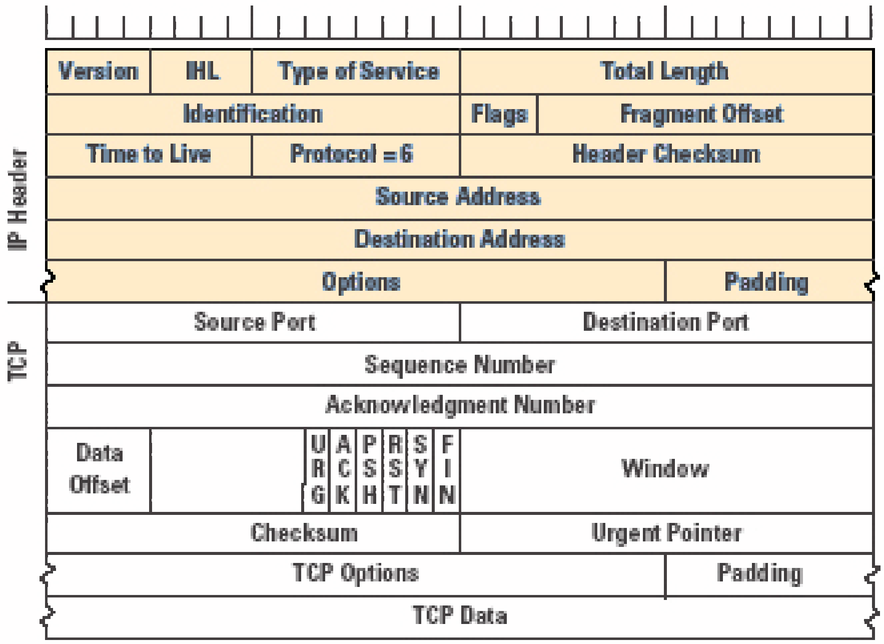

Figure 3-1: TCP/IP Packet Header

Although routers only need the IP header to forward packets, the entire TCP/IP header is sent "in the clear," or unencrypted, by default. This decision was not inconsequential, a point that was illustrated in the development of IPsec standards. IPsec encrypts the TCP header, a decision that was arrived at after considerable debate. The debate revealed the tension between revealing or hiding the TCP header from third-party observation. IPsec's eventual encryption of the TCP header emphasizes the purposes of each header: the router relies on IP for packet forwarding, but TCP headers are needed only at the endpoints.

The IP header includes, among other details, the source and destination IP addresses. The TCP header includes fields for source and destination *port numbers*. The port number is used by endpoints to manage system processes (or applications) that own networked data. A process sending or receiving data is assigned a port number, which enables the endpoint to match processes to packets. In other words, the port number enables demultiplexing of packets to flows of traffic matched to specific processes.

The IP and port numbers together identify a flow of traffic between two endpoints. The source and destination IP numbers identify the sending and receiving endpoints for a given flow.

The destination port number identifies the receiving process or application. The source port number is relied on to identify absolutely the flow and enable reassembly of the packets.[3]

Examples will help illustrate the uses of ports. A client connected to the Internet running an electronic mail (email) program, an instant messaging program (IM), and browsing the World Wide Web is managing three flows concurrently. These applications were assigned unique source port numbers when they opened a connection to the Internet. Each application connected to another endpoint's service (email, Web, IM) using that endpoint's matching destination port. Table 4.1 shows the request flows and response flows that might occur in this situation.

Requests from the Client				Responses to the Client			
Src IP	Dest IP	Src Port	Dest Port	Src IP	Dest IP	Src Port	Dest Port
128.30.5.43	18.980.3.14	51495	80	18.980.3.14	128.30.5.43	80	51495
128.30.5.43	220.3.43.233	51987	25	220.3.43.233	128.30.5.43	25	51987
128.30.5.43	18.2.33.120	52308	2980	18.2.33.120	128.30.5.43	2980	52308

Unlike clients, servers do not send requests to clients, but wait for incoming requests and respond appropriately. Clients must therefore know which of the 64,000 possible port numbers match the desired service on an endpoint. In initiating communication for the first time, clients rely on an external mechanism to learn the appropriate port. Today that mechanism is most commonly a list of globally well-known (port number, service) listings maintained by the Internet Assigned Numbers Authority (IANA).

The current architecture reserves port numbers 0-1023 for the most universal processes and applications such as Simple Mail Transfer Protocol (SMTP) or Hyper-Text Transfer Protocol (http). Ports 1024-49,151 are reserved for registered applications, those applications not standard or universal but that still use TCP/IP. An example of a registered application is AOL Instant Messenger, which matches to the registered port 5190. Port numbers higher than 49,152 are reserved for private or dynamic port numbers.

Servers listen for incoming packets on these reserved ports. For example Web servers listen for incoming packets on port 80, so browsers request information by sending a packet to the appropriate destination {IP address, port 80} pair. There is no technical enforcement of this mechanism; two parties could decide ahead of time to transmit Web traffic on any port so long as both parties knew which port. The list of well-known ports implies a convention used to coordinate initiation of communications.

This convention for port discovery is implemented in the application layer of the network rather than enforced by network architecture. Applications can and do neglect the use of well-known ports and even the list of well-known ports provides for applications that lack an assigned port. Port 1 is reserved for services to use specifically "instead of well-known ports."[50] Using port 1, multiple services can be accessed by sending a packet containing a service name. Positive responses will initiate connection using the specified service, negative responses indicate the service is not active at that IP address. Port 1 simply illustrates that well-known ports are not a necessity, simply a convention to which most application developers have agreed.

[3] While the entire four-tuple {source IP, destination IP, source port, destination port} is used to identify a flow, it is possible for two endpoints to share multiple flows for a given application at once: for example, a client running two web browsers to connect to http:web.mit.edu. In this case, the source IP, destination IP and destination port will all be the same. The two flows will *only* be differentiated by the source port number.

III. Network Address Translation

The development of NAT changed the way packets are forwarded through the network. NAT boxes act as routers by forwarding packets from the public Internet to a private network, but have capabilities beyond traditional routers. NAT boxes do not simply forward packets, but change header information when the translation from a global to a private IP address is made. Because NAT act as intelligent routers, some argue they break the end-to-end design of the Internet by adding functionality to the inside of the network.

Though there is more than one type of NAT, the most common one in use is the Network Address Port Translation (NAPT). The basic NAT simply translates the IP address; NAPT is a variation of NAT that changes not only the IP address but also the port number. Figure 3-2 shows the translation that occurs if a client behind a NAPT sends a packet to the Internet.

The interaction between port numbers and NATs influences the discussion in this thesis. Put simply, NATs extend the IP address space by borrowing bits from the port number.[72] This behavior undermines the purpose of using well-known port numbers.

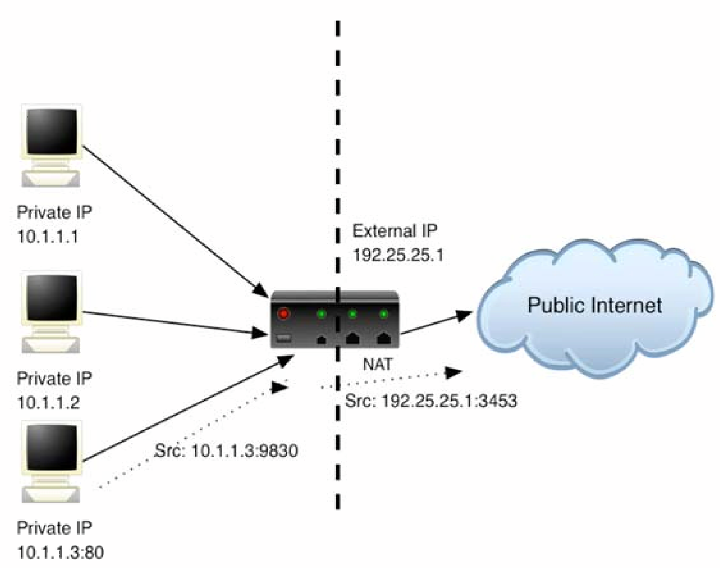

Figure 3-2 Network Address Translation (NAT)

As an illustration, take the situation where two web servers, A and B, behind a NAT box are listening for incoming packet s on port 80. To the public Internet, both servers seem to share an IP address, that of the NAT. When a host sends a packet to server A it is sent to the destination {NAT IP address, port 80}. Likewise, when a host sends a packet to server B it is sent to the destination {NAT IP address, port 80}. The NAT cannot be clear which server this packet is intended for unless the NAT has assigned port 80 to one of the two servers. If port 80 is mapped to server A, it is not possible for server B to publicize its services to the public Internet.

As this example illustrates, it is not possible to run two servers on the same port behind a single NAT. This has several implications, many of which were not accurately foreseen or fully understood in the early days of NAT implementation, among them the relative difficulty of running P2P applications and multiplayer games behind a NAT. In response, several means of circumventing NAT have been developed.

Two protocols are currently used to circumvent NAT: Simple Traversal of User Datagram Protocol Through NAT (STUN), and Traversal Using Relay NAT (TURN).[64][65] STUN and TURN are both client-server protocols, relying on an external server to enable machines behind a NAT to discover their public IP address and port number. STUN allows a machine to determine its public IP address and port number, as well as the type of NAT it is behind. This allows a machine to disclose its public address and available port. Neither STUN nor TURN allows for incoming requests from unknown IP addresses, but disclosing one's public address is useful for multimedia applications that may use a separate port for media transfer from the original well-known port used in discovery. For example, a machine using a P2P application will send an initial coordination request on one port, but receive data back on a second port. The originating machine must be able to tell others which second port to use.

A machine uses STUN to discover its own public-private IP and port mappings in order to tell a specific machine with which it is communicating what public address and port are open for receiving data. STUN is inadequate for the most restrictive NAT boxes, referred to as symmetric NATs, which change the mapping of private IP and port address to public IP and port address for each destination IP and port pair. In this case, the STUN server returns an address mapping that is only relevant for the private machine's connection to that server, not to any other machine on the network.

TURN is similar to STUN but also enables a machine behind the most restrictive NATs to have a single public address for all destination IP port pairs. A machine behind a NAT box opens a port by establishing communication with the TURN server. The TURN server acts as a relay by storing the address from which a TURN request is received, and assigning it a public address which routes to the TURN server itself. The machine then initiates communication with other machines and indicates that responses should be sent to its public IP as assigned by the TURN server. The TURN server receives these packets and forwards them to the address from which the original request was received. TURN is still considered work in progress and not an Internet standard; it is one among several tracks of research on enabling symmetric communication through NAT.[57]

IV. Summary: Key Port Number Characteristics

The port number is used today for two purposes, and corresponds to two types of identifiers. First, the port number partially identifies a flow of traffic, and specifically allows end nodes to demultiplex packets. Second, because well-known numbers are used, the port number identifies the type of communication associated with each packet, which enables identification of a flow of traffic as being Web, email, file-transfer or some other type of communication.

The first of these identifications is useful to the end nodes as well as third-parties by providing an aggregate identification for several packets as a single piece of communication. The second identification, type of flow, is arguably only useful for third-parties who do not have access to the end-system mapping of port number to process.

In the next chapter the consequences of using port numbers as dual identifiers of flow and type are examined. In partially identifying flows, the port number enables a key process, demultiplexing, without which packet-switched communication could not function. On the other hand, because it identifies the type of flow to observant parties, the port number enables second and third-order effects by providing a meaningful identifier to third-parties for traffic passing on the network.

Chapter 4 Diverse and Competing Interests

I. *Scoping the Tussle*

The tussle over revealing or hiding port numbers consists of the many diverse and sometimes competing interests of actors influencing and influenced by the Internet. The port number figured centrally in a recent FCC ruling on ISPs' ability to control customer VoIP usage, as well as ISP price stratification based on VPN usage. In those cases, there is conflict between users that would prefer hidden port numbers to prevent service blocking and ISPs that make business decisions assuming a revealed port number. Related to business decisions, ISPs have developed strategies for provisioning service based on the knowledge readily available because of port numbers.

The port number influences the ensuing security struggle between system owners and malicious attackers, between users and network administrators, and in some cases between governments and citizens. Firewalls that perform packet filtering are used to keep the "bad guys" out but also to monitor and control the "good guys" inside by inspecting port numbers. ISPs attempt to limit spam by using the port number to prevent customers from running mail servers. DOS attacks necessarily use port numbers to overwhelm services, and, accordingly, responses to DOS have been developed that rely on port numbers. Governments censor traffic and rely on revealed port numbers to ease the cost of doing so.

Application development has faced challenges because of standard use of well-known port numbers. Emerging applications that rely on a P2P model have to work around NAT boxes, which are problematic for P2P because of port translation. Application developers have developed interest in hiding the port number to enable use in the face of growing security and censorship controls.

This chapter seeks to elaborate on these tussles by evaluating each interest independently much in the fashion of a stakeholder analysis. The relevant actors include: ISPs, governments, users, enterprises, content owners, and application developers.

II. *Internet Service Providers*

Internet service providers (ISPs), particularly access service providers,[4] have matured in the architectural environment where port numbers are well known and sent in the clear. Today the port number reveals endpoint application identity to ISPs as traffic passes through their networks. ISPs have accordingly designed their business strategy around the port number, with the assumption that it will be possible to know which applications customers are running on the network. The port number is critical to several key ISP business strategies, including: service provisioning, traffic engineering, anti-spam measures, market segmentation and security.

A. Service Provisioning

Service provisioning is among the most critical business functions ISPs perform. Service provisioning involves predicting future demands on the network and allocating appropriate bandwidth to meet that demand. Internet service is provided in tiers; customers purchase service from third-tier access providers such as Verizon or Comcast, but often first- and second-tier backbone providers such as Qwest or MCI are responsible for the high capacity interconnected

[4] ISP is a general term that applies to entities all along the Internet service value chain. Access Service Provider refers to those ISPs that provide connectivity on the "last mile," directly to customers. For the purposes of this analysis, ISPs will almost always refer to access service providers. For further discussion see [4].

backbone. Access providers purchase bandwidth from backbone service providers at a usage-based fee, paying per bit transferred, but sell service to end-customers at a flat rate independent of actual usage. As a result, ISPs face an optimization problem: match uncertain demand with certain supply such that certain revenue offsets uncertain costs.

To solve this optimization problem, ISPs want to balance information asymmetry in their favor, specifically to reduce uncertainty in future demand. Monitoring traffic on the network reveals patterns of bandwidth use, and from those patterns future demand can be estimated. ISPs will find that demand varies with time-of-day, such that corporate customers demand higher bandwidth during the day and residential customers during the night. Demand also varies with type-of-use[5] so it becomes important to correlate type-of-use with time-of-day.[41]

Well-known port numbers reveal information about traffic that is used to reduce demand uncertainty. Port numbers reveal application type for a given flow of traffic which allows ISPs to easily discern patterns of use that might presage changes in demand or in traffic symmetry. For example, ISPs claim the information revealed by port numbers was crucial to provisioning in the face of increasing P2P use.[6] The introduction of P2P applications fundamentally shifted bandwidth symmetry and use. Rather than simply using downstream bandwidth to access central servers, residential end-customers began using higher upstream, and therefore more symmetric, bandwidth and larger quantities of total bandwidth. At first glance ISPs only see a pattern of increasing bandwidth use and symmetry.[41] To predict future demand ISPs want to understand what is causing the increased upstream use: is it an anomaly, or a pattern that is likely to continue?

In this case, the emerging pattern signaled a fundamental shift in residential network use. Residential customers, equipped with broadband connections, have become more active participants in the Internet's information exchange than just passive clients. If ISPs were unable to estimate and plan for the rising upstream demand that accompanied P2P, these changes in customer behavior would result in network congestion. The port number associated the majority of the upstream traffic with P2P applications, providing ISPs information that helped more accurately estimate future bandwidth requirements. Several emerging applications place similar demands on the network, most notably interactive multiplayer gaming.

As broadband deployment continues to expand, innovators and end-users will find more applications that use P2P network models and place similar demands on the ISP. VoIP use is one such application that is expected to grow, bolstered by broadband and competitive advantage in the telephony market. Currently, ISPs depend on the port number to identify these applications and the associated traffic to predict bandwidth demands and provision service appropriate to those demands.

B. Competitive Positioning

A central motivation of any firm in a competitive market is product differentiation. For ISPs this means differentiating service from that provided by competing ISPs to attract customers and establish market dominance. Port numbers are at the center of several mechanisms by which ISPs position themselves in a competitive telephony and Internet service market.

VoIP is an application through which ISPs have begun to differentiate their product and ensure competitive advantage. The introduction of VoIP fundamentally changes the competitive landscape of the telecommunications market. Because many ISPs offer traditional telephony service as well as Internet service, they have an opportunity to manage the demand shift from

[5] For example, P2P applications in general demand higher bandwidth than email.

[6] P2P refers both to a network model and an application. P2P network models are represented by relatively symmetric exchanges between endpoints. P2P most commonly refers, however, to the set of applications that have become popular for file-sharing, such as Kazaa and Grokster.

traditional voice to VoIP services by controlling service offerings. Entrant VoIP service providers, such as Vonage and Skype, use the physical infrastructure provided by incumbent ISPs to offer phone service, but compete with the incumbents in telephony service.

Incumbent ISPs want to limit their customer's ability to use third-party VoIP services using their infrastructure. Instead ISPs would like to force customers to use their telephony services and switch to their VoIP service at the ISP's pace. ISPs do not want to carry third-party VoIP traffic if they could collect revenue by capturing those customers themselves. There have recently been a handful of instances in which ISPs have been reported to be blocking traffic to the VoIP ports, preventing customers from using VoIP service offered by Vonage, Skype, and similar providers.[27]

This type of port blocking creates competitive advantage for the ISP. Those companies in control of network traffic have the advantage to block third-party services, forcing subscription to their own services. This behavior could also evolve into product differentiation between ISPs: ISPs might market themselves as "the ISP that allows you to use third party VoIP service."

The FCC issued a ruling prohibiting ISPs from blocking the ports commonly used for VoIP services in March, 2005.[28] The move does not preempt ISPs from using the same competitive mechanism on other applications, only with VoIP. The policy introduces several interesting questions, including uncertainty about which applications will be deemed protected by government regulation and what justifies that protection in which cases. It also opens the possibility that ISPs will find other means to control application use in competitive positioning.

Without the ability to block VoIP ports, ISPs may use other technical mechanisms to ensure competitive advantage over third party VoIP providers. For example, ISPs might simply degrade quality of service to third-party VoIP applications, behavior that might be difficult to track. Quality of service cannot easily be tested because it is dependent on traffic congestion. Furthermore, because the Internet is a best-effort service, ISPs might find immunity in the inability to ever guarantee quality of service.

Today the ability to control application use on an ISP's network depends on well-known port numbers to identify application type. More costly analysis, such as deep-packet or timing analyses, could reveal similar information. ISPs favor the well-known port number as an easy, readiably available and meaningful identifier to ease service differentiation.

C. Market Segmentation

In addition to gaining market dominance, ISPs are interested in capturing the most revenue possible. The common business practice is to segment the market based on willingness to pay, allowing a company to charge higher prices to those customers who would pay them while retaining a broad customer base. The knowledge gained by monitoring port numbers allows ISPs to segment their market. Two examples illustrate how ISPs use port numbers to segment the market: VPN applications and email traffic.

One of the earliest charges of port-based market segmentation was the blocking of port 550, the port used for VPN. With an increase in telecommuters required by corporate policy to use VPN, ISPs recognized an opportunity to demand a higher price for the service. Widespread accusations emerged that various ISPs were blocking VPN.[35] The solution: order a business package from your ISP. Most business packages offer a static IP address but are otherwise no different from residential packages, unless ISPs decide to block ports for residential customers as they did with VPN.

Most recently ISPs have segmented the market by blocking port 25, which also enables spam control as will be described below.[44][49][8] Though ISPs justify blocking port 25 with the spam concern, many savvy customers with broadband want to run their own mail servers, which is impossible if port 25 is blocked.[7] In the face of consumer outrage ISPs might have just opened port 25, but instead leveraged the situation to increase revenue through price differentiation.

Many customers simply do not notice port 25 blocking; most residential users don't need to run a mail server or access servers other than that provided by the ISP. Only the savviest of users are negatively impacted by port 25 blocking, and, as documented by use-net discussions across the Internet, these users tend to call their ISP with complaints. In response, the ISP chose to cater *only* to the demands of these users, offering them access to port 25 if they were willing to pay a business subscription fee.

These two examples are the most widely recognized and documented examples of price differentiation and market segmentation based on the port number, but it is not difficult to imagine emerging market segments based on P2P, VoIP, or interactive gaming applications. While some will argue that this price differentiation is not a business necessity, it is nonetheless a behavior ISPs have come to rely on and expect the ability to carry out. It is also a benefit ISPs enjoy as a corollary to port blocking motivated for any other reason, as demonstrated by the spam example documented below.

D. Spam Reduction

Spam is the digital version of junk mail: unsolicited email. Because the marginal cost of sending email is zero, so-called spammers have no incentive to limit or restrict the amount of spam sent. As a result, inboxes are inundated with spam and network resources are unreasonably consumed. ISPs want to restrict spam to free congestion of network resources and provide better service to customers. To this end, ISPs install spam filters on their mail server to filter spam out from meaningful email. They also rely on ports to restrict spam proliferation.

Simple Mail Transfer Protocol (SMTP), illustrated in figure 4-1, runs on well-known port 25 to send outgoing, unauthenticated email. A client machine sends packets to destination port 25, and mail servers listen for incoming packets to that port. ISPs often run their own mail server, as do enterprises, educational institutions, and even individual users. A Verizon subscriber might send SMTP packets to port 25 on the Verizon mail server, but may also wish to use SMTP with the mail server at his company or school. He may also want to run his own server, effectively sending SMTP to port 25 on his own machine.

[7] It is possible to run a mail server on some other port, provided all users accessing that server know which port it is on.

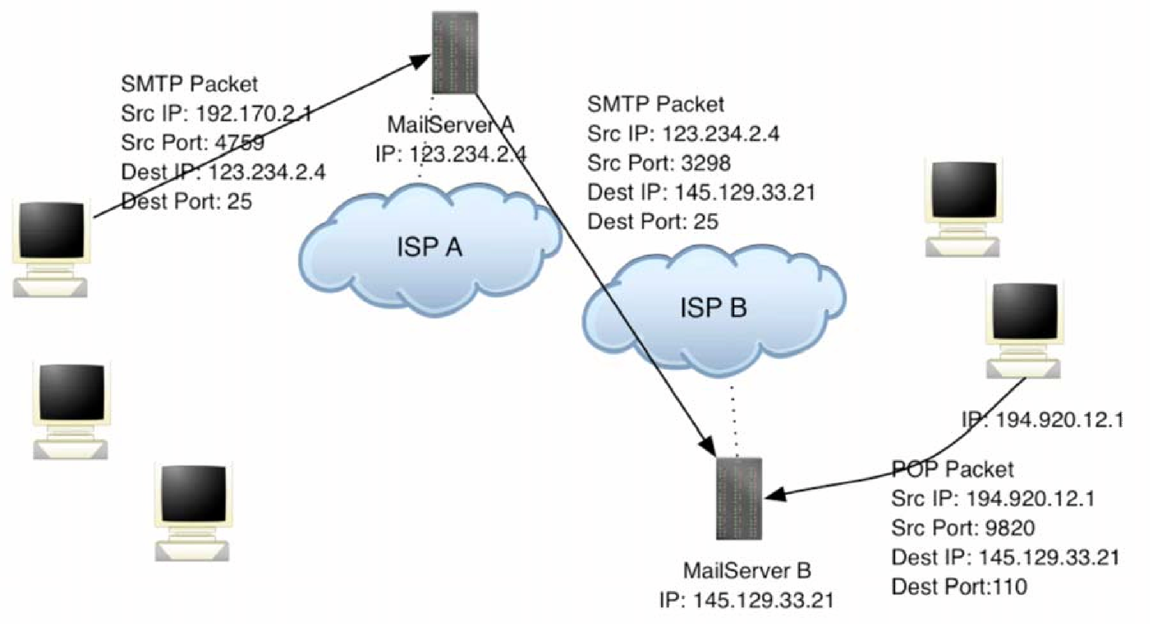

Figure 4-1 SMTP

Spammers distribute spam through Trojan programs and viruses to avoid accountability and distribute processing requirements. Once installed on an unwitting consumer's machine, these Trojans are programmed to run SMTP and send email packets to destination port 25 on some machine, hijacking it as a mail server.[73] This mechanism thwarts most efforts by ISPs to control spam by filtering their own mail server.

In 2004 ISPs began blocking outgoing packets to destination port 25 other than those sent to the ISPs mail server IP address.[7] This effectively limits a customer's ability to access outgoing mail servers other than that run by their ISP. Customers are unable to run their own SMTP mail servers or access external SMTP mail servers. But spam is limited somewhat because the viruses and Trojans installed by spammers to send SMTP mail via port 25 are also unable to do so.

Blocking port 25 could be an example of an ISP business practice masquerading as a socially beneficial policy. In fact it is an effective means of controlling spam, but ISPs have also taken advantage of the opportunity to segment the market as described earlier. To open port 25 one must only upgrade from residential to business service.[8] Of course, SMTP is not the only protocol by which mail is sent: users can easily send authenticated mail or use SSL without interference by the ISP.[8]

E. Security & Risk Mitigation

Port numbers identify network traffic destined for vulnerable applications. ISPs, network administrators and individual users install firewalls to filter traffic destined for vulnerable ports as a means of limiting viruses and other attacks. Firewalls control connections to and from an internal network and log port scans. While port scans are not attacks on a system or network, they are used to identify ports actively listening for incoming packets and often precede an attack.

ISPs and network administrators use the port number to mitigate risk to the overall network posed by vulnerable applications. In most cases market segmentation occurs as a result of this behavior: ports that are blocked for security reasons are opened for business customers, most likely with the expectation that network administrators actively monitor and secure vulnerable

[8] Authenticated mail and SSL require authentication that cannot be achieved easily by viruses and so are not susceptible to spam hijacking.

ports and applications. Similarly, network administrators will often open ports for specific users if a legitimate request is made, but otherwise block those that correspond to vulnerable applications.

A well-known example of security motivated action is the blocking of incoming connections to port 80. Port 80 is the well-known port hyper-text transfer protocol (HTTP), or more commonly the Web. In 2001 two viruses garnered wide-spread attention as incredibly destructive, both because of the speed of their diffusion and the amount of damage caused.[9] Code Red propagates itself by initiating connections to port 80 on random hosts and, when a host is listening for incoming port 80 requests, Code Red exploits HTTP commands to eventually deface web content. Nimda spreads through email initially, but once it has infected a web server will use port 80 to spread to users accessing the server.[11][12]

To combat both these worms, and several less well-known ones, ISPs began blocking incoming connections to port 80 for residential customers.[9] Customers can still access external web servers, but cannot run their own. For many customers, there is no need to run a web server; for those who want access to port 80, they can request access from their ISP at the cost of a business subscription.[10]

Several other ports are regularly blocked to residential customers for security purposes. These include, most notably, 135, 137-139, 445 and 593.[63] These ports translate to a set of applications that use parts of the Windows operating system with well-documented security vulnerabilities.[73] It is rare, if not unheard of, for an ISP to block ports on business-class customers, despite these security vulnerabilities.[11] It follows that any security-reasoned port blocking serves the dual actions of mitigating security risks and segmenting the market.

Port blocking in the name of security is risk management. Because the majority of residential customers are relatively unaware of the details of network security, ISPs bear the security risks for these customers and the secondary risk to their own networks. Blocking ports is one way ISPs manage this risk. For those customers that want full functionality, the ISP is willing to award it so long as the customer is willing to offset the security risk by paying a higher fee.

III. Government

The government's interests reflect the diverse interests of other actors in the market, be they individual citizens or private firms. For simplicity, government interests can broadly be divided into two sub-sections: the interests of law-enforcement and of market regulation. These must be treated separately to acknowledge the distinct political, technical and regulatory interests associated with each. Though the focus of this thesis is in the United States, to avoid analyzing a topic of global concern in an overly narrow context, examples of government stakeholders will be taken from current political governance structures throughout the world.

A. Law Enforcement

Governments generally have an interest in enforcing the laws of the state, and the Internet can broadly be seen as a relatively new communication medium on which crimes can be committed and evidence of wrongdoing collected. As such, the law enforcement arm of government may have an interest in monitoring communications across that medium or limiting behavior on it. Some countries have decided to monitor and filter all Internet traffic, while others limit the amount of government monitoring that takes place. Regardless of where on this spectrum a given

[9] Code Red uses an infected machine to deface web pages. Nimda modifies web content and several executable files on infected machines. Both worms have the potential to cause economic damage to firms that rely on web servers for business purposes.

[10] It is common now for ISP user agreements to include a restriction on customer's use of the service to run external, or public, servers of any kind - email, Web, or otherwise. See for example [20]

[11] Often network administrators will block those same ports for an enterprise's private network.

government falls, it is likely all governments depend to some extent on well-known port numbers to reveal user behavior.

China has erected a nation-wide firewall to filter Internet traffic entering, leaving and within the country. It is not the only nation known for Internet censorship; Cuba, Laos, North Korea, Saudi Arabia, Tunisia, Syria and Vietnam are also recognized as nations that censor Internet traffic.[74] In most cases this involves filtering web traffic based on content, not just the application in use. For cost-effective content-based filtering of web traffic, a firewall is best-served by knowing with which applications incoming packets are associated, and making a first-level filter based on that information. If authoritarian regimes were faced with performing deep packet inspection on each and every packet, censorship would be far more costly.

The United States requires that specific criteria be met before Internet traffic can be monitored in real time.[1] Most notably, when law enforcement agencies gain legal authority to wiretap an individual's communications in real time, they are required to meet a "minimization" requirement: to tap only relevant communications and ignore irrelevant communications that may occur over the line during the course of a wiretap.[12]

Knowing the applications in use allows this "minimization" to be done in a cost-effective way: by focusing only on those applications of interest, be they email, P2P, or web. Controversy over VoIP surveillance provides an example of the tradeoffs at stake. The Federal Bureau of Investigation has demonstrated interest in extending CALEA requirements to ISPs as a way of guaranteeing access to VoIP communications.[13][51] It is unclear how CALEA compliance will be implemented technically, however it is likely that well-known port numbers will play a key role in identifying packets associated with the application of interest.

B. Market Regulation

Government regulators are concerned with encouraging an economic environment in which innovation can flourish and in which critical infrastructure providers, in this case ISPs, are able to develop sustainable business models. Regulators act as facilitators of diverse interests, reflecting the interests of ISPs, application developers and end users at once. It is near impossible to be entirely balanced, and under any given administration, market regulators make decisions that influence various actors more than others. Well-known port numbers are one additional variable in the Internet space on which regulators can exert power.

Until recently, the FCC had not applied regulatory power to packet headers. As mentioned earlier, the FCC in March 2005 declared that ISPs cannot block VoIP ports.[27] FCC Chairman Powell indicated the decision was in part intended to ensure the "Internet remain an open platform for innovation."[28] This attempt to promote innovation depends entirely on there being well-known port numbers. Of course, the violating action, port blocking, depends on well-known port numbers as well.

A second area, outside of pure regulation, in which governments might influence the market is Internet taxation. Though it has yet to be implemented, it is not difficult to imagine a government that decides a tax on email, on VoIP, or on P2P use would be reasonable. Such taxes could be implemented today because each of those applications is uniquely identified in packet headers via the port number.

[12] Minimization does not apply to stored communications, only those intercepted in transit. To date many Internet-related wiretaps have accessed stored communications in email, but consideration for real-time surveillance is warranted as well.

[13] Communications Assistance for Law Enforcement Act requires that ***telecommunications*** providers implement technical capabilities that enable law enforcement to easily and readily access information when a warrant is served. As of yet, these requirements have not been extended to information service providers.

A third way in which port numbers figure into regulators' actions is the government's interest in maintaining a secure infrastructure and one that is relatively free of spam. As discussed in detail under the ISP section, well-known port numbers figure centrally into security and spam control. In the United States the government does not directly act to improve security or reduce spam for the public Internet, but through the United States Computer Emergency Readiness Team does issue guidelines on security for the private sector and to state and local governments.[75] Governments such as Saudi Arabia and China, on the other hand, are likely to take a much more active role in monitoring network security and spam. All these actions currently depend on the use of well-known port numbers to filter traffic to vulnerable applications and to control spam.

IV. Hardware Manufacturers

In general backbone hardware does not need the port number to perform its primary function of forwarding and routing packets, all the necessary information is provided in the IP header. There are some notable exceptions where hardware depends on the port number for functionality. NAT boxes rely on the port number to do port forwarding. While the port translation enabled by NAT boxes has arguably enlarged the IP space by allocating multiple machines to a single IP address, it can be problematic in preventing multiple machines from running the same service behind a single NAT.[69]

Firewalls, often though not always implemented in hardware, depend on well-known ports to do packet filtering.[13] While firewalls can be configured to block traffic from specific IP addresses, they are most often configured to generally manage network service connections. Firewalls can block services altogether or only from specific hosts based on well-known port numbers. Firewalls also are able to indicate when an IP address has been the subject of a port scan, often a sign of an upcoming attack.

Many hardware providers do make routers and switches that have capabilities beyond the simple packet forwarding function. Functionality is added to support stateful packet inspection, in other words content inspection beyond the headers.[13][29] Other routers enable Quality of Service (QoS) implementation by inspecting the Type of Service (TOS) bits in the packet header, which can be used by applications to indicate if there are low latency or low throttle requirements on the traffic.[3][34] Neither of these capabilities depends on well-known port numbers.[14]

The backbone manufacturers build products to meet the demands of the ISPs, whether those are simply packet forwarding, firewalls, or QoS enablers. Their products provide ISPs with the capability to identify network traffic as associated with applications and perform analysis on that traffic to meet business requirements. To an extent, the interests of backbone manufacturers mirror those of ISPs: they want to meet the demands of their customers. On the other hand, adding complexity to the hardware, such as deep packet analysis, would enable these manufacturers to raise their prices substantially and increase revenue.

V. Application Developers

Any application designed to run on the network relies on the port number to enable coordination amongst independent machines. Well-known port numbers make it straightforward for servers to listen for requests from unrecognized clients: any packet arriving for port 80 is automatically recognized, regardless of the source port, as a packet intended for the web server. From this initial communication a conversation can be developed between two endpoints on the network.

[14] The port number and TOS bits were intentionally separated in TCP to prevent linking application type to quality of service received.

Application developers are increasingly moving away from the traditional client-server model of Internet communication to P2P models. This shift changes the role of the port number in coordinating networked applications, particularly when NAT is taken into consideration. The following application examples provide some insight to how emerging networked applications are using well-known port numbers in the NAT-enabled environment. What is perhaps most important is to note the barriers to innovation that existed for these applications, because the unknown innovations are the most uncertain and therefore hardest to predict.

A. VoIP

Many VoIP technologies use the Session Initiation Protocol (SIP) to initiate communications between computers. SIP is a generic session protocol, but has been used most widely in deployment of VoIP. SIP uses a well-known port, 5060, at the source and destination to initiate a stream of media (voice) which is then passed over a second, dynamically and randomly assigned port.[37]

To run VoIP using SIP, a machine should be listening on port 5060 for incoming packets. This can be problematic if multiple users want to run SIP behind a single NAT box. The common solution has been to use SIP proxy servers that keep track of all logged in users; this way users can send packets to the proxy servers, from which the public IP address and port number for machines behind a NAT can be determined using protocols like STUN or TURN.[64][65] At some point in the network, a VoIP-enabled router is needed to translate from IP to the traditional PBX network. These can be placed at the endpoints or, as is often the case, at the edge of an enterprise network.[15]

As has been discussed above, VoIP has been the subject of ISP port blocking activity in recent months. Independent VoIP service providers were at the mercy of ISPs who may have chosen to block port 5060 at any time to secure competitive positioning or simply manage traffic. Clearly this behavior was not in the favor of third-party VoIP developers.

This changed in the United States in March, 2005 when the FCC issued a Consent Decree with Madison River Telephone Company, LLC, a small ISP in Virginia that had become the subject of an investigation on VoIP port blocking.[27] The Consent Decree makes it illegal for ISPs to block ports designated for VoIP traffic. This regulatory action seems to favor independent VoIP providers, but third party VoIP traffic is still subject to inadvertent packet loss or poor quality of service, actions ISPs may take next in place of port blocking to preserve their competitive advantage.

Blocking or degrading third-party VoIP traffic illustrates one of the challenges application innovators face. Because incumbent ISPs offer traditional telephony service they have had little incentive to spur growth or development of VoIP service. Arguably VoIP's development has occurred because of start-up innovative companies wiling to take a risk on the technology and large enterprises that demand the cost reductions in telephony fees. If third party VoIP vendors are subject to port blocking and service degradation, the reward for their innovation is captured by incumbent ISPs that may have played very small roles in the research and development.

B. P2P Applications

P2P applications face the same challenges in working behind NAT boxes as VoIP: it is difficult for more than one host to run the same P2P protocol behind a single NAT. P2P is also constrained, as is VoIP, by the problem of self-discovery: an endpoint cannot easily advertise its service to the public Internet if it does not know which public IP address and port number correspond to itself. Protocols like STUN and TURN have eased these challenges somewhat.[65][64]

Most P2P networks today use "supernodes" to facilitate rendezvous between users. These supernodes can be any user on the network, but are generally those with the highest bandwidth and fastest hardware and must have a public IP address (i.e. cannot be behind a NAT). Supernodes collect the identifying information, IP address and port number, from other users and broadcast it on the network. As such, the supernodes act as central servers, allowing those users behind NAT boxes to publicize their public IP address and port number without knowing it themselves.[43]

P2P lacks the industry support enjoyed by VoIP because thus far it has only gained publicity for being a source of illegal sharing of copyrighted material. Court battles over Napster and most recently Grokster cast P2P technology as a enabler of copyright violation.[2][39] The questionable legal status of these applications positions developers and network administrators contentiously.

Because P2P applications have been widely used for illegal file-sharing, many network administrators have tried to control their use on internal networks. The most straightforward way to do so was initially to block the well-known ports associated with the most common P2P applications: Kazaa, Morpheus, BearShare, FastTrack and Grokster.[76] In response, users and developers moved the applications away from well-known P2P ports to benign well-known ports, most often port 80 (HTTP), a behavior known as port tunneling. Port tunneling masks P2P requests, making them look like standard web requests and thus benign.[26]

Most recently, application developers have designed the software to use "port hopping" to avoid detection.[67] Port hopping refers to an application's frequent movement amongst random port numbers. Port hopping makes any attempt at port blocking futile, and has led to the development of deep-packet inspection and recognition of application-signatures beyond the port number.[41][38] The next step seems to be encryption of P2P traffic, which makes identification by network administrators and ISPs incredibly difficult.[41]

C. Other Applications

P2P and VoIP offer insight into the challenges application developers face in deploying applications based on the P2P model in the context of well-known ports and NAT. Other applications, most notably interactive gaming applications, have faced the same challenges and taken similar steps to overcome them. Gaming applications enjoy relatively little backlash from network administrators because they are perceived as a more legitimate use of the network than P2P applications. Nonetheless, the technical challenge of working around NAT is universal for any P2P modeled application and represent an interesting tussle between users who might be indifferent towards revealing the port number, but hide it in order to simply overcome the NAT barrier.

VI. Intellectual Property Owners

The interest of content owners are relevant to this discussion because of the sheer financial and political power they have been able to assert in the Internet space, most notably around P2P applications. Because content is simply carried across the network via these other applications - gaming, P2P, etc. - the port number does not play a direct role in production or ownership of content. However, well-known port numbers allow network administrators and ISPs to monitor and track application behavior such as P2P, and therefore give them reason to control and monitor those applications.

The Recording Industry Association of America (RIAA) began filing lawsuits in 2003 against individual users accused of copying and sharing copyrighted music recordings. Though some lawsuits were found to be in error, the majority were filed based on the IP addresses of users the RIAA deemed "egregious" violators.[45] The RIAA was getting IP addresses from inside the P2P networks, rather than monitoring all traffic behavior across a network. Content owners simply want users to reveal their identity to P2P services to enforce accountability for illegal actions. The port number is entirely irrelevant to the content owner's direct actions.

Content owners have pursued legal recourse against P2P as a technology that does no more than enable copyright violation. If these efforts fail, it is likely that content owners will continue pressuring ISPs and network administrators to cap down on P2P use. Content owners will be helpless bystanders, however, as the arms race exacerbates between P2P developers concealing application-identifying information and ISPs performing deeper inspection techniques to identify applications.

VII. Enterprises

Enterprises are in some ways end customers of network resources, but primarily they are intermediary agents of control. Enterprises are interested in maintaining the security of their private network, controlling spam, and monitoring independent user behavior inside the network in order to enforce network policy. To this end, enterprises deploy firewalls and monitor traffic much in the way ISPs do.

Firewalls protect a private network in several ways, most notably through port-based packet filtering and often by acting as NAT boxes. Preventing the public Internet from directly accessing machines on an internal network by using NAT may complicate running some networked applications, but it also prevents attacks on those machines by masking the private IP addresses.[72][69]

Firewalls commonly filter incoming packets based on port numbers. It is widespread practice to prevent traffic from accessing ports 135-137 because the associated applications are known to be vulnerable to attack.[73] Port-based packet filtering can be used to bolster security, as shown by that example, or simply to enforce network policy. If users agree that when accessing the network they will not run instant-messaging or P2P applications, the ports associated with those applications can be blocked as a mechanism of enforcement.

Network administrators also block outgoing packets based on port numbers to assert control over use of internal network resources. An enterprise might block all outgoing traffic except that sent to the enterprise web proxy, through which the enterprise can easily monitor and control interactions with the public network.[40] Though more detailed packet inspection is required to identify transmission of pornography, an action commonly prohibited on enterprise networks, this inspection is aided by well-known port numbers in isolating web traffic.

Finally enterprises are often subject to DOS attacks which can cost enormous financial damage. DOS attacks rely on well-known ports to identify target services before inundating those services with packets.[52] Responses to DOS attacks rely on port numbers less frequently, but instead rely on stateful packet inspection to identify attack signatures deeper in packet content. Recent research suggests that dynamic ports may be useful in mitigating DOS attacks, specifically addressing the applicability of port hopping to the problem.[46] This may suggest that well-known ports make systems more vulnerable to DOS attacks.

VIII. End Users

End users have possibly the most complex set of interests of any stakeholder discussed. They want it all: the ability to run applications freely behind NAT boxes and without interference by ISPs, as well as strong network security and freedom from the deluge of spam. Additionally, user-friendly interfaces to the network are favored.

The primary conflict users are involved in is between controlling entities, network administrators and ISPs, and users seeking free use of the network. Users in the United States are often shocked and horrified to learn that an ISP has been blocking ports as revealed by discussions on various technology-related forums.[70][76][8] This outrage illustrates the cultural mindset of Americans regarding Internet use: access, once subscribed to and paid for, should be open and free without restriction. Many users feel that individuals should deploy firewalls, similar to how an enterprise would as described above, rather than being subject to ISP-enforced port-blocking policies. Clearly, this cultural attitude conflicts with ISPs' interests in mitigating risks and protecting competitive advantage.

Arguably this outrage is not entirely cultural. In China, where communist has bred a culture of censorship, efforts to subvert the government-imposed national firewall abound. Several software applications have been developed for the purpose of subverting government censorship in China, and some, such as Freenet, have succeeded thus far. In other words, the battle between users who seek freedom and entities that control the network is likely to rage in any setting.

IX. Summary: The Arms Race

The emergent behavior of this system is akin to an arms race. The major players have divided themselves into two camps: those interested in hiding the port number to preserve anonymous behavior on the Internet and those interested in revealing it to aid monitoring behavior on the Internet. In the first camp sit ISPs, law enforcement and enterprises. In the second, application developers, end users and arguably regulators. The weapons include port tunneling, port hopping, and eventually encryption on one side. On the other the primary weapons are port blocking and deeper packet inspection enabled by faster processors and cheaper memory.

The arms race is well underway in the battles being fought in the P2P, VoIP and VPN spaces, and will likely spill over to most networked applications. The first action was port blocking to limit and control use of certain applications: P2P, SMTP, VPN. In response, users and then application developers began using unknown ports and subsequently port-tunneling over "benign" well-known ports once deep packet inspection got underway. When port-tunneling wasn't sufficient enough evasive behavior, those seeking anonymity began using port-hopping as a means of obscuring not only the application, but the stream of traffic into pieces.

The system has finally reached a level of complexity and conflict that regulators have entered the fray. Though at first glance the FCC's decision to disallow port blocking on VoIP seems like a win for VoIP and innovators, its implications are uncertain and complex.

Will the ISPs counteract by degrading third-party VoIP service? If they did, it seems likely the FCC would attempt to regulate this behavior as well; they clearly have an interest in protecting VoIP providers' access to the infrastructure. It is, however, unclear how such regulation would work mechanistically. If VoIP service is degraded, it's possible the market would evolve to favor ISPs with "open access" for third-party VoIP providers. However, given the relatively concentrated nature of telecommunications services and the necessary nature of its use, consumers have limited choice of ISPs and exert relatively little market power. Might VoIP providers have an interest in devising a technical solution similar to port hopping to evade detection by ISPs? Possibly, or they might perpetuate the arms race to an "endgame" status: fully encrypted packets.

A bigger question is how and when the FCC decides to regulate port blocking for other applications. This regulatory action could lead Internet innovation to hinge as much on lobbying power in Washington as technical sophistication or market dominance. In other words, if applications are only "protected" from ISP port-blocking once the FCC steps in, application developers will most certainly want their fair say at FCC hearings, and those most successful in Washington will gain access to the network.

The way the conflict over access to the port number is progressing, a likely end-state to prepare for is end-to-end encryption of traffic. When faced with deep-packet inspection, users and application developers will have incentive to encrypt packets entirely to hide application identity. In the current system, there is no incentive structure for users to do otherwise, but end-to-end encryption would endanger the ISP business model. Few would argue it is a desirable end-state. This arms race for access to the port number poses a challenge to engineers building future architectures to think about alternative designs that might influence or simply mitigate the conflict.

The next chapters will diverge from the status-quo analysis offered here, and focus on uncertain consequences of using random as opposed to well-known port numbers. As described in this chapter, actors rely heavily on the port number because it is linked to a well-known list of application types. It will be useful then to think about the opposite state of affairs, random port numbers, where the sixteen bit number is only useful at the endpoint and carries no meaning in the middle of the network.

Chapter 5 Implementing Random Port Numbers

Earlier I introduced the port number as it fits into the current technical context of Internet architecture. As discussed, today most Internet applications operate with the use of well-known port numbers,[15] a somewhat arbitrary design decision that has resulted in the cat-and-mouse game between those interested in hiding application identities and those interested in revealing that information as discussed in the previous chapter. Looking backwards on the decision to use well-known port numbers and its effects leads us to wonder how things might have been different. The question is, if it were possible to enable greater choice between inadvertent and intentional revelation of application identity, how would this impact the conflicting shared and individual interests on the network?

In this chapter I will briefly introduce several research proposals that suggest a movement away from well-known port numbers. In many ways these proposals have not reached widespread or full implementation yet, and I am not advocating any single idea. The point here is to explore the basic requirements on the port number as well as creative ideas for its flexible uses.

I. Context for Technical Change

It is necessary at the outset to explain the context of these research proposals. We have in a sense reached a state of technical ossification with the Internet. The hardware on which this virtual world is built is widespread, standard and will be difficult to change unless the benefits vastly outweigh what would be immense financial costs associated with uprooting the physical infrastructure. And though it would be far less costly to change the software and protocols on which the Internet runs, such a change still calls for vast reconfiguration of an incredibly complex communications infrastructure.

Any proposals to change well-understood and accepted standards should therefore be met with initial skepticism: how possible will it really be to implement these changes? For the purposes of this analysis, a leap of faith is necessary. Make no attempt to divine how we might get to this new world, simply consider what might happen were we to snap our fingers and communicate on an Internet in which one of these or other unknown proposals, in more mature form, was the standard. The following chapter will take this analysis a step further and explore "what might be," but first the basic technical requirements and flexibilities warrant revisiting.

The port number as used in TCP/IP is simply a demultiplexing tool. In concert with the IP source and destination numbers, the source and destination ports identify a flow of traffic. This is the most basic requirement on the port number, that it signals to end machines of which flow each packet is a part. The decision to make these numbers well-known added a second feature that is more flexible: the port number is now used in discovery, i.e. to coordinate communication with a service. When requesting information from a Web server, the initial request is sent to port 80. To move away from well-known ports, it will be necessary to design a mechanism that will coordinate initial communication with a service.

It warrants comment that using port 80 for Web traffic is enforced at the application level today. The application designers in a sense “own” the coordination mechanism, though savvy users acting in concert are able to hijack ownership of the mechanism for their own purposes. For the most part, however, application designers dictate and control the coordination mechanism for networked applications by conforming to the convention of well-known ports. As demonstrated

[15] With a few exceptions such as a handful of P2P applications that have implemented random-ports as a counter-measure to port blocking, most applications on the Internet today are identified by a well-known port number as maintained by IANA.

by some file sharing applications today, this need not be the case. Random ports could be used, today's convention could be abandoned altogether. Standard mechanisms, however, would facilitate new conventions to which application developers might subscribe. Alternatively, a new network architecture could be designed such that the coordination mechanism mimicked that used for IP coordination, the DNS. Such a mechanism would be owned not by application developers but ownership would be decentralized throughout the network.

The following research proposals suggest ways of initiating communication beyond the current use of well-known port numbers, allowing standard widespread use of random port numbers. The three proposals elaborated on here are examples of how random ports might be implemented at varying layers of the architecture. First, port hopping represents and example of how application developers might incorporate random ports into software. Second, DNS SRV examines how random port numbers might be implemented in overlay networks. And finally, the FARA architecture is presented as a prescription for building an architecture that fully incorporates random port numbers.

II. Implementations

A. Port 1

The first implementation of random ports was in some sense enabled in the very early days of the Internet. Port 1, as described earlier, was reserved among the list of well-known ports as the single port on which services that did not use a well known could be multiplexed. Requests sent to port 1 include a string describing the desired service. Responses can either be positive, '+', or negative, '-'. A positive response initiates communication using the specified service while a negative response indicates the service is not present. Although it is a simple implementation of random ports, port 1 does not seem to be not widely used.[50]

B. Port Hopping

To begin understanding additional ways of how random ports might be implemented it is logical to start with what has already been implemented today. Port hopping, which is simply dynamic port assignment, has been implemented today by some P2P applications in order to evade identification by network administrators and firewalls. The number can be changed multiple times and frequently, making identification even of a single flow more difficult than simply tracking the four-tuple created by source and destination IP and port numbers.

While port hopping is one instance of using random as opposed to well-known ports to evade identification, it is not a full implementation of random ports. Those applications that use port hopping often use well-known ports in the discovery process before switching to dynamically assigned ports.[10] If, on the first attempt, connection fails, the application attempts to connect to another randomly selected port picked from a pool of pre-established ports, sometimes including ports generally assigned to other applications (such as Web on port 80). Once the initial connection is made, P2P clients move to random ports. Using this mechanism, P2P has evaded detection by most network administrators and successfully circumvented many attempts to block its use.[16]

C. DNS SRV for Service Locations

RFC 2782, a proposed Internet standard, describes a mechanism for locating services without using well-known ports.[36] This proposal is the most mature among current research in this area,

[16] P2P has evaded detection through port numbers by using the port hopping technique. Other forms of detection that look deeper into the packet or at traffic patterns are constantly under development.

but has yet to be widely implemented or used. Essentially, it suggests adding a level of indirection to the DNS such that services could be looked up in the same way locations are now. In much the same way that the DNS maps a string to an IP number, DNS SRV maps a string to a port number.

Moving port numbers into the DNS means that the mapping of service to port number can be updated more frequently and easily than is done today, allowing use of dynamic port numbers. Unlike port hopping, which does not enable discovery, DNS SRV uses the DNS mechanism to enable service discovery. It is therefore a far more complete implementation of random port numbers than done by port hopping.

D. FARA

FARA is a general framework for organizing architectural components. It describes a way of thinking about network architecture by bringing together a set of ideas that had been independently proposed into one central organization. FARA separates the demultiplexing requirement from service identification, allowing for random numbers to be used in demultiplexing and using a rendezvous server to coordinate service requests.[18]

In many ways the DNS SRV fits this aspect of FARA, though because FARA describes a general framework for an entire architecture DNS SRV is not an implementation of FARA. Nonetheless, the idea of separating the two current meanings of ports numbers (demultiplexing tool and service identification) is one that is advanced by the DNS SRV as well as FARA.

III. Summary: Coordination Requirements

The most basic requirement of implementing random ports is a reliance on a coordination mechanism for service discovery. Port hopping relies on well-known ports as the initial mechanism while DNS SRV relies on the DNS for its coordination. FARA assumes some form of coordination service without specifying what that will look like, offering the possibility of a look-up service or process listings within a larger system.

The implementations of random ports that have been deployed most widely to date, notably port hopping and DNS SRV records, represent two distinct models of coordination. The first is compatible with today's model of application ownership over the coordination mechanism; port hopping is implemented at the application layer. The second represents a departure from this model by decentralizing ownership of the coordination mechanism, a departure also suggested by the framework outlined in FARA.

Ownership of the coordination mechanism is critical because this mechanism is the control point for implementing random port numbers. Without it, it is impossible to know on which ports a communication can take place. The degree of flexibility in the mechanism as well as the level of centralization in its ownership will influence how fully a random port number system can be implemented.

Chapter 6 Reactions to Random Port Numbers

Imagine it was possible to immediately change the Internet architecture. Instead of using well-known TCP port numbers, the sixteen bit number would be negotiated between two hosts on the network such that every conversation was represented by a pair of random numbers. Examples of how, mechanistically, this might be implemented were introduced in the previous chapter without attention to the steps required to change from the current implementation to a different one. Here, we are interested in examining the consequences of this change rather than technical details of how we might get from today's network to one that relies on random ports. For the purposes of this analysis, then, a leap of faith is necessary.

Imagine a world in which port numbers are random rather than well known. Such a design might be desirable for a number of reasons, primarily to mitigate or dissolve the conflicts being played out today. In short, random port numbers would empower users and restrict third parties' ability to interpret users' actions on the network. Based on the earlier analysis a slew of questions arise about how ISPs, application developers, enterprises or governments might react. Similarly, random ports enhance security in many ways but also have the potential to derail many mechanisms relied on today. Coordination for P2P modeled applications will be enhanced with random port numbers, but network administrators might lose control of enterprise traffic.

This chapter seeks to understand how each of the actors might respond to random port numbers by anticipating actions and reactions based on how business and application needs have evolved to date.

I. Internet Service Providers

Anticipating the moves and counter-moves under a hypothetical architecture begins with the ISPs. ISPs are arguably the dominant stakeholder in this system: their power is derived from ownership of the physical infrastructure, their financial might, and their role as service providers. Based on the dominance of ISPs, it is logical to begin this analysis with expected ISP responses to random port numbers. From there, counter-moves by other parties - including users, application developers, content owners and governments - can be strategically anticipated.

A. Business Requirements

If the architecture shifted to random port numbers, ISPs would have to adjust their business strategies for service provisioning, service differentiation, and market segmentation. All these strategies today depend on well-known port numbers to identify application use in network traffic as previously discussed. Though any anticipation of responses to an architectural change are speculative, the exercise of thinking through strategic possibilities is a necessary prerequisite to design changes.

Service provisioning is among the most fundamental business functions ISPs perform: it is necessary to prevent congestion on the network. Without application-level knowledge of packets traveling across their networks, ISPs will be unable to accurately identify patterns of application-use that signal impending shifts in traffic symmetry and bandwidth demand. A system of random port numbers, whether implemented at the application layer or deeper in the network protocols, would undermine current service provisioning strategies.[17] ISPs will have to change current

[17] Random port numbers might be implemented by a single application, as has been seen in some P2P protocols aiming to subvert identification. It can also be implemented broadly by changing the underlying network protocols. See Chapter 5 for a discussion of how ISPs currently rely on well-known port numbers for service provisioning.

pricing models to create incentives for users to provide application level information to the network, or find alternative ways of accessing that information.

There are of course other ways of inferring or deducing the type of application using information external from the port number which, although less helpful for market segmentation or service differentiation, could be useful in service provisioning. These include inferences based on destination IP addresses; connections to IP addresses known to be associated with domains hosting major Web servers such as cnn.com are likely to use distinct applications from those directly between users, for example.[18] The size and length of the connection also reveal something about the application in use, and can be particularly useful in statistical aggregation.[38] It is also possible in some cases to identify applications based on characteristic signatures.[68] Further development of any of these techniques might be explored by ISPs.

The dual business goals of service differentiation and market segmentation would not be achievable via today's common technique, port blocking, if random ports were implemented. Segmenting one's market and differentiating one's product, in this case Internet service, are both critical to any firm in a competitive market. Without the ability to block certain applications for some customers and not others, or to control third party applications running on the network, ISPs may face commoditization. ISPs would likely seek to develop innovative pricing models to enable market segmentation and service differentiation.

B. Potential Responses

To maintain a competitive market with random port numbers, ISPs will necessarily shift their pricing models to reflect the lack of information available about user's application usage. It is well accepted that customers greatly prefer simple, flat-rated pricing models, as opposed to usage-based models.[56] Based on this, it is unlikely ISPs would institute usage-based pricing then, even though doing so would allow for straightforward market segmentation and ease of provisioning. It is more likely that ISPs will fundamentally change what they sell as opposed to how they sell it: Internet service as a function of bandwidth alone is difficult to support financially without the type of information currently provided by well-known port numbers.

As an example of how ISPs might fundamentally change what they sell, ISPs could begin selling incoming and outgoing port numbers rather than simply IP addresses. This would essentially shift ISP property rights to include services, as opposed to simply bandwidth per machine. ISPs could implement this pricing strategy today, it is not one that is technically enabled by random ports but simply one that might be more attractive to ISPs if random ports were standard.

While this strategy would de facto segment the market and provide an opportunity for service differentiation, provisioning would remain difficult. A user might purchase four incoming port numbers that he randomly assigns numbers to for four different applications, but his ISP cannot know whether those applications are asymmetric or symmetric, high-bandwidth or low-bandwidth, bursty or real-time. One strategy might be to offer unlimited outgoing ports, but charge a premium on incoming ports to control the bandwidth symmetry. Additionally, the ISP could impose bandwidth limits on incoming ports to control incoming traffic.[19]

A second pricing strategy that might appeal using random ports would be one that has a rich academic history but has yet to gain industry support: a quality of service (QoS) pricing model. QoS has been lauded as a necessary mechanism for supporting the wide variety of applications that run on the network today.[22] Several technical implementations have been proposed,

[18] Note that the IP address alone will not reveal the domain associated with it, and with random ports it would not be possible to intercept DNS traffic. Still, an ISP might know from external sources that a given IP number hosts the CNN Web server, and from this might infer the type of traffic sent to that destination.

[19] This technique would mimic that used by many web hosting services. See [78] for example.

including IntServ, DiffServe and RSVP. Though end customers can purchase routers that support these QoS mechanisms at the endpoint, ISPs have yet to implement QoS support in the core of the network because there is currently no business case to make the investment in hardware upgrades.[5][24]

Random port numbers may create a business case for QoS because a QoS-based pricing model would provide ISPs with information about user's demand for service that is today assumed by knowing the requirements of applications in high use. Whereas today well-known port numbers reveal applications in use that have latency demands, for example, a QoS pricing mechanism would force users to reveal this information to the ISP without relying on application identification but instead relying on the Type of Service (TOS) bits.[20]

By presenting users with a set of QoS options, the market would automatically be segmented and product differentiation would naturally occur. High priced QoS would simply reflect a user's willingness to pay for transmission of packets with high TOS bits. VoIP, for example, might work best if TOS bits represent low-latency and low-jitter requirements. ISPs could price these requirements such that loss to third party VoIP applications would be offset by the ability to charge higher rates for the type of service required by those applications. P2P, on the other hand, might not have high service requirements and users might set TOS bits to appropriately represent a low priority for these packets.[21] As a result, P2P would only consume high bandwidth when network congestion was such that other network traffic would be unaffected.

By prioritizing available bandwidth based on network congestion, a QoS pricing model enables service provisioning despite a lack of detailed application layer information.[22] Today provisioning is a practice of matching uncertain demand with certain supply. QoS is a mechanism for distributing certain capacity according to certain willingness to pay as demonstrated by subscription to a QoS price. In other words, QoS would solve the provisioning problem created by random port numbers by managing uncertainty in user demand.

An additional business opportunity may arise for ISPs as a result of random port numbers. Any implementation of random port numbers at the architectural level, as opposed to the application level, will require an intermediary port coordination mechanism as discussed in the previous chapter. ISPs might be able to leverage this requirement by offering an additional service to customers in the form of coordination. In other words, ISPs could attempt to leverage their power to create a market for reliable coordination mechanisms and charge users accordingly.

C. Security & Risk Mitigation

Current practices to mitigate security are based on using firewalls to block well-known port numbers that have not been “opened” by a machine behind the firewall. Ports are blocked unless specifically opened by a service. If random port numbers became standard, ISPs current practices of blocking incoming ports would not be effective in accomplishing the goal of mitigating network risk. New approaches would be required.

Current approaches to limiting spam by blocking incoming traffic on port 25 will not be possible with random ports. Instead we could imagine mail servers, especially those run by ISPs, agreeing to a set of conventions to govern mail transfer between servers. These conventions might mimic well-known ports, requiring incoming mail transfers to use an agreed-upon port such as port 25, but would function in an architecture that supported random ports. Such conventions would in some ways preserve the status quo by allowing ISPs to block outgoing port

[20] TOS bits are in the packet header similar to port numbers, but are intended to indicate the type of service required for each packet. The TOS bits are the inputs to QoS implementations.

[21] P2P, unlike VoIP, will work sufficiently well if packets have high latency or high jitter or are subject to congestion and arrive over a long period of time.

[22] Service provisioning will require QoS models that support latency, jitter and have a bandwidth limit.

25 as they do today. However, mutually agreeing users could opt-out of the convention, theoretically making themselves vulnerable to spammers, but simultaneously allowing individuals to run their own mail servers. In effect, such a convention would create an opt-out mechanism from conventions enforced by social norms.

Port scans, the common precursor to malicious attacks, will be less meaningful if not meaningless altogether with random ports. Without well-known ports, the range of ports expands from its current range of approximately 1,024 that map to system applications to a range of 64,000 that might map to system applications. The sheer number of ports that would require scanning may make port scans computationally costly, and obviously identifies malicious behavior.[23] Additionally, an open port does not reveal the service active on it, so attackers would be presented with two challenges. First, find an open port. Second, find the service active on that port and determine its vulnerabilities.

If port scans did become useless, firewalls may also lose some of their value. In addition to logging port scans, firewalls perform varying degrees of packet filtering. Today, a lot of packet filtering is based on the port number alone; this type of packet filtering will not be broadly useful, but only if mechanisms allow users to inform the firewall of specific ports in use that require traffic monitoring. Firewalls have additional value beyond packet filtering, though, including prevention of address spoofing and blocking specific IP addresses or domains. There are additionally application-level firewalls that understand specific application protocols and headers, free from relying on port numbers, and these will remain useful with random port numbers.

To regain the kind of control ISPs currently exercise on user's ability to receive incoming requests they would need to leverage port coordination mechanisms. It is entirely possible that ISPs might offer coordination services to users, but it would be difficult, though not out of the question, for them to leverage the mechanism to control incoming requests. To do so, an ISP would have to closely link the coordination mechanism with routers at the edge of its network and monitor all outgoing traffic in real time. Given the complexity of maintaining such a close real-time link to the edge hardware, it's unlikely to be a cost-effective strategy for ISPs generally.

II. Government

A. Law Enforcement

The law enforcement arm of government interests is likely to initially be frustrated and stifled by random ports. Whether the government in question aims to control content on the Web or simply gain access to voice communications of suspected criminals, law enforcement organizations are generally interested in knowing as much as possible about behavior on the network. While ISPs are interested in application use as a means of estimating other information (bandwidth requirements) governments are interested simply in application use that is not readily available without well-known port numbers.

An obvious step the government could take would be to inspect packets beyond the headers, what's known as deep packet inspection. Doing so reveals the full content of the packet, often making it possible to identify not simply the application in use but also the content of the communication. Deep packet inspection is required to filter Web content, for example; well-known port numbers reveal that a packet is http and deep packet inspection is used to filter the actual content of that Web traffic. This illustrates an important point: deep packet inspection is possible, but not used widely because it is not cost-effective. In other words, well-known port numbers today allow use of high-cost analysis on select packets. With random port numbers, the

[23] Any port scan identifies the likelihood of impending malicious attack, but onr of the magnitude required to scan 64,000 ports would be blatant to any intrusion detection system, making maliciously painfully obvious for attackers to undertake.

high-cost analysis would be necessary on *all* packets simply to find the select packets of interest, making its cost prohibitive if many users are targeted, as is the case in censorship. However, law enforcement is often interested in the communications of a select group of targeted users, and in these cases the analysis is not so costly as to be prohibitive under random ports.

More interesting than surveillance is the possibility of government interference with port coordination mechanisms as a means of control. The coordination mechanism could be used either for censorship or prohibition of illegal activities. Whatever the implementation, random port numbers will require a coordination mechanism that, as identified earlier, could conceivably be controlled by the ISP or third parties. Because these mechanisms will be the loci of control for application identity, they will be targeted by law enforcement agencies seeking to monitor communications. For example a government might seek ways to install "wiretaps" at a coordination mechanism to inspect requests from a suspect's IP address, accessing the assigned port numbers for applications of interest such as email.

In the US a clear application of interest is VoIP, which law enforcement would like to treat similarly as it does traditional telephony.[51] In addition to a wide array of other technical challenges facing law enforcement's ability to monitor VoIP, with random port numbers law enforcement will lose the ability to easily identify packets as VoIP. A response to random port numbers in this case is hard to predict, especially given the already complex technical issues facing VoIP monitoring today.

B. Market Regulation

Random port numbers seems to be in direct conflict with the recent move by the FCC to establish regulatory power over ISPs support or restriction of certain applications by limiting port blocking. The FCC is unlikely to oppose random ports, however; it's response is likely to be neutral because their objective, ensuring user's ability to run third-party VoIP applications, is furthered by random ports. Regulators will be interested in how ISPs respond, however, particularly because the FCC is interested in facilitating network investment that requires a strong ISP business model. But the FCC has not yet regulated Internet service pricing, as opposed to traditional telephony service, which is subject to pricing regulation.[21]

Governments interested in applying taxes to specific applications will be challenged to identify those applications. Just as law enforcement will face a cost-prohibitive deep packet inspection on each packet, so would any attempt to implement application-specific taxation via traffic monitoring. The result may be Internet access taxation, similar to taxes that are currently applied to local exchange carriers whose infrastructure supports telephony and many ISPs. Tax burden could be shifted to end-users via general usage-based taxes, which may be met with some opposition.

The traditional approach to government-run network security is undermined by well-known ports for the same reasons ISPs will no longer be able to use today's mechanisms for security and spam control. The government's response is likely to be regulatory. In some nations the government may seek complete control of all available coordination mechanisms, allowing control over services. This is likely to be difficult. With random port numbers, a government would have to keep a log of all "conversations" approved in its coordination mechanism and block all unapproved packets, similar to what ISPs could do with control of the mechanism. Again, this would require real-time linking between the coordination mechanism and hardware but for those nations that place a high enough value on control, it would be technologically possible.

Of course, technically this would not in itself prevent the use of third-party coordination mechanisms. To do that, a government would have to verify all services it registered and ensure none were simply third party coordination mechanisms. This may result in a virtual chamber of commerce: to operate services in this nation, you must agree to the following policies and register with our port coordination mechanism.

A second regulatory approach to security might push security regulation further out from the center of the network, mirroring the shift in technical capabilities. . With random port numbers, network security is necessarily pushed further out than in today's system. Packet-filters no longer act as security firewalls: it is not possible to simply block incoming packets to vulnerable applications before they reach the endpoint. A regulatory approach might shift the economic and legal liability for security vulnerabilities to the endpoints themselves, either to applications developers or end users. Holding application developers liable for security has been promoted in the past but never implemented; random port numbers may change the technology enough to force the market to adapt.

Finally, many governments are interested in fostering innovation as a means of economic growth. Arguably, random port numbers, by preventing third parties from interfering with emerging applications, are more likely to encourage innovative applications. This depends, however, on there being a variety of available coordination sites. Regulation may become necessary to prevent conglomeration of coordination sites, but is likely to be no different than standard antitrust policy as exercised in other markets.

III. Application Developers

By many accounts application developers stand to gain with random port numbers, as illustrated perhaps most tellingly by the number of networked applications today that have moved to using random port numbers despite conventions that favor well-known numbers. However, in some cases random port numbers may complicate an application's ability to gain wide support. ISP reactions to random port numbers will significantly affect application deployment, as well.

Most notably, application developers will face a simplified task in designing peer-to-peer applications to run behind NAT boxes. The current difficulties associated with NAT arise from its inability to map multiple private machines to a single publicly well-known port number. With random port numbers, this mapping will no longer be necessary. Instead, different private machines will be allocated different, random external port numbers for the same services.[24] By easing the task application developers face in servicing from behind a NAT box, random port numbers should allow a greater variety of applications to emerge.

A. VoIP

VoIP is an example of a real-time application on the network that cannot function well with the traditional best-effort service or bursty nature of Internet traffic. Though random ports will free VoIP providers from traffic blocking or degradation by ISPs and will ease VoIP functionality from behind NAT boxes, the real interesting implications for VoIP occur when potential ISP reactions is considered.

In recent years VoIP has gained traction because of emerging hardware and software solutions that prioritize VoIP packets at the edges of the network. Though the ability to prioritize VoIP has been critical in deploying VoIP on congested private networks, VoIP packets are not currently prioritized on the backbone of the public Internet. Because there is not a business case for introducing QoS into the core of the network yet, VoIP gets the same quality of service across the backbone as all other data. This works fine today because the backbone is over-provisioned and at the bottlenecks, the edge of enterprise networks, TOS bits are used to prioritize VoIP packets.[14]

[24] It is worth noting that random port numbers are not necessarily more effective as a mechanism to get around NAT boxes than TURN.

VoIP providers would be interested to see ISPs introduce a QoS pricing model. If such a model existed, users desiring VoIP service would receive the necessary QoS as demonstrated by their willingness to pay for it. This is likely to be a good thing for VoIP deployment, but has the potential to instigate counter responses from the government in the arena of universal service. In particular, the government may shift its focus from universal PSTN service to universal Internet service.

In the US access to PSTN telephony infrastructure has been the focus of universal service policy. In the case of universal Internet service, the growth of VoIP may shift the focus from universal access to infrastructure to access to service.[33] The universal service debate with respect to Internet service rages on, and is likely to grow more complicated with growth of VoIP which effectively merges the debates of telephone and Internet universal service. If ISPs begin offering service on a QoS pricing scheme, it is possible the FCC will step in to regulate Internet access prices based on the same logic used to justify their regulation of telephony access today.

B. P2P

P2P developers have already recognized the benefits for their applications of using random port numbers. Random port numbers have been used in P2P as a strategy to avoid port blocking imposed by network administrators. The games between ISPs and P2P today illustrate the conflicts that are likely to emerge and have informed much of the analysis in this chapter. The conflicts in this space have evolved into a game for control over information, with several P2P applications beginning to use encryption as the final action to protect information.

P2P will also respond to how ISPs change Internet pricing schemes. The QoS pricing model has been the focus of this analysis so far; in analyzing P2P the drawbacks of QoS pricing emerge. As initially suggested, a QoS pricing model would allow users to choose from varying brackets of QoS. These brackets might vary on any number of QoS variables, latency or jitter for example, and will also include some measure of available throughput. P2P networks will likely thrive or struggle based on the aggregate membership's willingness to pay, which may lead to stratification in P2P networks.

An alternative approach ISPs might try to accommodate the variety of QoS demands any one user may have would be to sell QoS per port number. An analogous model would be wireless subscription services in the US. Users subscribe to plans that offer a given number of minutes per month during peak hours and unlimited calling in off-peak time, while minutes beyond the allotted amount are charged additionally, usually at a high price/minute. These plans distribute limited cellular service according to a willingness to pay for possible permeations of several variables.[30]

Similarly, selling service as some function of QoS and port numbers would distribute a scarce resource, bandwidth, among multiple applications according to willingness to pay. A user might opt for a plan that offers high QoS for one incoming port number, which he then chooses to use primarily for VoIP, several other incoming port numbers at medium and/or low QoS, and unlimited outgoing port numbers. This pricing model would empower the user to choose what service he desires, while also giving ISPs necessary feedback that reveals user preferences and service requirements.

C. Unknowns

The most interesting question to think about, but perhaps the most difficult, is how random port numbers will affect emerging, as of now unknown, applications. With P2P, VoIP and other well understood applications, the change to random port numbers influences a tussle space we can identify based on behavior today. As these applications developed, their technical capabilities influenced the tussle space they fell into, and as a result fully understanding that tussle space beforehand would have resulted in what today would be perceived as a limited understanding.

Similarly, understanding the tussles that might be of interest in the future and provide a context for emerging technologies are hard to understand today.

Unknown applications will present a dilemma for ISPs and others interested in understanding network use. Should ISPs continue current efforts to identify traffic through pattern analysis, selective deep packet inspection, etc., what will happen when they see traffic that is unidentifiable, representing a new application? Possibly the ISP would be able to recognize certain properties of a new application and identify it repeatedly as "application X."[38][68] This form of pattern matching might be sufficient for ISPs' purposes, but knowing more, an application's purpose for example, may require research or knowledge beyond what is accessible on the network.

IV. Content Owners

Content owners will find themselves more powerless to stop the advance of P2P than they do today should random port numbers be implemented. Their reactions are likely to be similar in many ways to those seen today in response to P2P applications that have already begun to use random port numbers. For these applications, the port number has lost its meaning as an application identifier and in response other mechanisms are developed to identify P2P.[68] Content owners may not be directly involved in the development of these mechanisms, but certainly their interests motivate research in identification of application signatures in network traffic.

Content owners will be affected most directly by their inability to pressure network administrators into limiting P2P and illegal file sharing via port blocking. Given that they have several avenues by which to prevent copyright violations, they may simply have to accept an inability to influence P2P use, especially if the current legal uncertainty falls on the side of P2P developers.[25] With random port numbers, content owners will have to address illegal file transfers directly rather than influence overall P2P use. Approaches to do so have been tried already. The RIAA and other media organizations have infiltrated popular P2P networks with bogus files and spam, for example, to create negative incentives for file sharing.

An alternative approach has been modeled after the tactics used against ticket scalpers or drug dealers, both of which are similar in many ways to copyright violators. Illegal transactions occur between two willing parties, and go unnoticed so long as the two parties are able to prevent third parties from monitoring the transaction. It is therefore difficult to hold the scalper or dealer accountable. Similar to undercover police, content owners have played the part of a willing party to illegal transactions and directly sued for copyright infringement. In the absence of well-known port numbers this may become more prevalent behavior directed not simply at egregious suppliers but also the common user downloading a handful of files.

V. Enterprises

The enterprise reaction to random port numbers will provide further illumination of the security implications of the architecture change. The most fundamental purpose of port numbers to the enterprise is as a mechanism of control: traffic is screened, filtered or denied based on well-known port numbers today. In essence, the well-known port number has enabled a security approach that is very attacker focused. Liability for security vulnerabilities under this system falls to the network administrators and attackers.

Random port numbers may first and foremost shift the focus of security liability. If it is no longer straightforward to apply traffic filtering based on applications known to be vulnerable,

[25] The outcome of the *Grokster* case will have a significant impact on content owners' avenues to stop illegal filesharing.

enterprises may seek methods and policies to focus liability inside the enterprise at end nodes. A possible outcome may be individual employees held liable for security vulnerabilities opened by applications they run. Company policy may forbid employees from running P2P because the applications are known to introduce viruses. If John Doe runs a P2P application inside the enterprise and introduces a virus, John Doe could be held accountable for damages inflicted on the enterprise.

An interesting outcome is one that has been proposed in academic literature but has yet to gain traction in the legal world: application developers become liable for damages inflicted because of security vulnerabilities. This could result if enterprises feel powerless to defend against malicious traffic directed at vulnerable applications. Without a method of defense, enterprises may seek to shift the risk to external entities, most likely to application developers.

VI. End Users

End users are at the center of this system, affected by many of the actions and reactions of others but powerless individually to affect the system in meaningful ways. Users will be strongly affected by changes to ISP service models, and in the aggregate will have consumer power to influence which pricing models are most successful. Consumers have demonstrated interest pricing systems where the cost is clear up front, favoring flat rate over usage rate models.[56] This preference can be seen in wireless plans as well: bucket plans in the US allow most users to avoid per-minute charges, and likewise in Europe it is common to buy minutes up front.

A service and QoS variable pricing model might illustrate a consumer preference for flexibility, as well. ISPs and telephony companies today offer plans that consumers can choose for long term, amendable at relatively long intervals such as months or years. Consumers will likely want the ability to change their service plan at shorter intervals if it includes any variation on number of port numbers.[26]

In reaction to the likely move by content owners to seek retribution against even the most commonplace copyright transgressors, users are likely to seek P2P solutions that integrate file encryption into the system. An example of P2P software that is likely to gain widespread support under this system is Grouper. Grouper was built by a group of friends who wanted to share family videos with each other easily, but avoid strangers from accessing those files. It is a P2P network that fully encrypts all data that is shared between users, and, as an added benefit, allows users to connect only to small groups of subscribers. The result is somewhat similar to social networks such as Friendster; users are able to share files with a group of trusted friends and acquaintances, fully protected by the strength of encryption.

Grouper has gone to great lengths to protect itself from attack by RIAA, specifying in its documentation that files can be streamed or downloaded but that, like FreeNet, the central application has no knowledge of which files are shared. The system also makes it impossible to share files with strangers and limits the size of a group. By limiting the size of the group, the software is approaching the realm of fair use more closely than traditional P2P applications have in the past.

An interesting push and pull will emerge around the coordination mechanism. As the locus of control, the coordination mechanism is the technical piece around which users stand the lose the most should control be unilaterally imposed by ISPs or governments. When user's freedom to deploy applications has been limited by ISPs or governments in the past, it is rarely a corporation or application developer that initiates counter-reactions. More often, individual users find a

[26] On signing up a user may use three applications frequently: email, Web, and VoIP. The day he discovers a new application, such as an innovative P2P solution, he will want the flexibility to upgrade his service immediately.

mechanism to subvert the controlling authority which, in successful cases, gathers steam and develops into a widely used technique or application.

The question will be what types of subversions will be possible if coordination mechanisms are monopolized. One can imagine "shadow services" that emerge to cover up a second coordination mechanism; it looks like a user is connecting to a legitimate service, when in fact he is just getting through the first coordination mechanism to get to another. It is also likely that very tech-savvy users will try to find mechanisms to demultiplex multiple applications over one port, essentially adding a degree of indirection to TCP.

VII. Summary : Shifting the Power Balance

Random port numbers, above all else, would realign information symmetry between individual users and large ISPs, governments and even network administrators. It is not entirely clear that the end-state of end-to-end encryption would be avoided, but there is clearly a possibility of reaching a different outcome.

If random port numbers were implemented, ISPs would lose a piece of information that has, to date, been central in traffic management and service provisioning. The first likely shift would be from ISPs attempting to recover that information, by deeper packet inspection, using application signatures, or monitoring timing patterns. Such a shift is unlikely to necessitate a dramatic response on the part of consumers or application developers.

The second likely shift is that ISPs will attempt to recover the ability to segment their market and differentiate their services. The mechanisms mentioned above may work for implementing general business policies to an extent, but will not be as effective as the well-known port number in absolutely identifying applications subject to screening. Instead ISPs would probably alter pricing strategies as discussed to ensure that service prices better reflect customer's willingness to pay.

Another interesting shift would appear in network security. Firewalls as implemented today would not be very effective with random port numbers, but it is unclear if attackers would use port scans because they would be easily detected, and certainly attackers would be unable to target vulnerable services based on the associated well-known port number. A distributed port scan might ease the computational load and would be difficult for firewalls to respond to effectively, but an open port would not automatically identify the service running on it. The inability to defend against networked attacks would thus be mirrored by a similar inability to mount networked attacks. The result might be a shift in security focus from the edges of the network directly to the endpoints and the associated software. Alternatively the coordination mechanism could become a target, if fooled it would provide access to vulnerable machines.

Finally, random port numbers are most likely to empower users to run applications and services with greater freedom than today. This point relates to the ISP's increased difficulty in blocking services or applications, but also is related to the enterprise who faces a similar difficulty. With random ports, it would be more difficult for enterprises to monitor and control user behavior. As a result both consumers and employees would be empowered to run programs of their choice.

Chapter 7 Conclusion

At just sixteen bits, the port number has had a far-reaching impact on the business structures, uses and vulnerabilities of the Internet. Though used first and foremost as a way to identify flows of traffic to endpoints on the network, the well-known port number has been used for a wide set of capabilities across the network. In particular, it has become a central piece of the architecture as it relates to security and control, business development and innovation.

Reevaluating the unintended consequences of designing port numbers to be well known will inform decision makers responsible for designing applications or entire architectures in the future. In this process, it is helpful to consider alternative design choices and evaluate how things might be different. By juxtaposing the well-known and random port number designs, this thesis set out to not only inform designers, but also to provide an early example of how the uncertain consequences of architecture design can be evaluated.

Unfortunately, beyond analyzing random and well-known port numbers separately, this work does not provide an objective recommendation in port number design. Any attempt to do so inherently incorporates values and becomes a subjective recommendation. Nonetheless, from the analysis and presentation of data, it is possible to make some statements about the benefits and consequences of each system.

Most notably the arguments around port numbers deal represent a battle for information. The well-known port number identifies more than a flow, it also identifies user behavior in such a way that network administrators and ISPs have an interest in the information. In a number of cases, users would prefer this information remain hidden while network operators depend on accessing it. It seems that random port numbers would tilt the tussle space in favor of users, while well-known port numbers tilt it toward network operators. Such a tilt cannot, however, be measured. ISPs arguably rely on the well-known port number because they can, not out of necessity. In that case, random port numbers may balance the tussle rather than tip it.

While ISPs might be able to design innovative business models and provisioning strategies to accommodate random port numbers, it's not clear that other capabilities could be as easily adapted. For example, the security landscape would be forever altered without well-known port numbers. Today's firewalls and the attackers they defend against would become less relevant to network defense. Spam and security control will be pushed to the end-user and away from network operators, and while users could always opt-in to a network-operated security system, it's unclear how ISPs or network administrators would regain that control network-wide. It’s possible this shift would motivate greater use of intrusion detection systems, which notify systems of anomalous behavior that might signal an attack. This result would signify a more adaptive security model than the current reliance on well-known port numbers for service identification.

The port number is also inextricably linked to NAT and therefore entrenched in the arguments about innovation. Arguably the current combination of well-known ports and NAT creates a barrier to innovative P2P modeled applications. Applications like P2P and VoIP have had to overcome the difficulty of running services behind NAT created by port translation. The characteristics of random port numbers are such that port translation presents less of a barrier.

An architecture that supported random port numbers would realign the information imbalance between end-users and third parties in such a way as to neutralize it. Today, under well-known port numbers, third parties have easy access to more information about a user than vice versa. In turn, users are subject to control by third party network operators and security systems depend on an unreliable identifier.

The tussle has evolved such that random ports are already being implemented at the application level already, and the consequences are being played out in real time. ISPs are developing other tools to identify applications for the cases when well-known ports are no longer used. The question may not be if an alternative design would be beneficial, but what to expect as the alternative arrives.

It is important to note the difference between application ownership of the coordination mechanism and decentralized control. If applications continue to move to random ports independently, the coordination mechanism will be unique for each application. This may be desirable in that it distributes vulnerability; if one coordination mechanism were hijacked somehow, others would still function. On the other hand, application ownership of the coordination mechanism means that the transition to random port numbers may be slow, delaying if not preventing benefits from being accrued to the system.

Should the current push towards application ownership of random ports continue, we can expect a mixture of ISPs' evolving reactions to random port numbers to include deeper packet inspection, and security systems' inability to adapt quickly to this new model may create vulnerabilities on the network. These increased vulnerabilities may trigger more controlling responses on the part of firewall administrators, potentially creating barriers to innovation. A full architectural shift would be required for smooth transition, and advisable in spite of anticipated opposition on the part of some stakeholders. It seems likely we are moving to random ports and away from well-known ports in any case, and a concerted decision to do so as a system will ease the transition.

Chapter 8 References

[1] 18 USC § 2518(5). U.S. Code, Procedure for Interception of Wire, Oral or Electronic Communications.

[2] *A&M Records et. al. v. Napster, Inc.*, 284 F.3d 109 (9th Cir. 2002).

[3] Almquist, P. "Type of Service in the Internet Protocol Suite." RFC 1349. July, 1992.

[4] Altmann, Jörn. "A Reference Model of Internet Service Provider Businesses." ICTEC2000, 3rd International Conference on Telecommunication and Electronic Commerce, Dallas, Texas, USA, November 2000. Accessed at: http://www.i-u.de/schools/altmann/publications/publications.html 20 March 2005.

[5] Bell, Gregory. "Failure to Thrive: QoS and the Culture of Operational Networking." *Proceedings of the ACM SIGCOMM 2003 Workshops*. August, 2003 in Karlsruhe, Germany.

[6] Blumenthal, Marjory and D.D. Clark. "Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New World." *ACM Transactions on Internet Trechnology.* Vol. 1, No. 1, August 2001, pp 70-109.

[7] Bode, Karl. "Blocking Port 25 Traffic: 'MyDoom' Virus Reheats Discussion." January 29, 2004. BroadbandReports, Accessed at http://www.broadbandreports.com/shownews/38004 April, 2005.

[8] BroadbandReports.com Forums, SBC West-Ameritech, "Port 25 Party Poopping." http://www2.dslreports.com/forum/remark,12480346~mode=flat

[9] BroadbandReports.com Forums, Up and Running, Technology Law and Politics, "ISP blocking port 80…why they doing this?" Accessed at http://www.broadbandreports.com/forum/remark,4846921 April, 2005.

[10] Caviglione, Luca. "The 'Dark Side' and the 'Force' of the Peer-to-Peer Computing Saga." *P2P Journal.* January, 2004. Accessed at http://p2pjournal.com/issues/Jan04.pdf April, 2005.

[11] CERT Advisory CA-2001-23 Continued Threat of 'Code Red' Worm. July 26, 2001. Accessed at http://www.cert.org/advisories/CA-2001-23.html April, 2005.

[12] CERT Advisory CA-2001-26 Nimda Worm. September 18, 2001. Accessed at http://www.cert.org/advisories/CA-2001-26.html April, 2005.

[13] Cheswick, William, S. Bellovin, and A. Rubin. (2003) *Firewalls and Internet Security.* Addison-Wesley. United States of America.

[14] Cisco Systems. "DiffServe – *The* Scalable *End-to-End* QoS Model." Accessed at http://www.cisco.com/warp.public/ccc/pd/iosw/ioft/iofwft/prodlit/difse_wp.htm April, 2005.

[15] Cisco Systems. "Security in SIP-Based Networks." (2002) Sipsc_wp.pdf. Accessed at www.cisco.com/warp/public/ cc/techno/tyvdve/sip/prodlit/sipsc_wp.pdf April, 2005.

[16] Clark, David D. "Name, Addresses, Ports, and Routes." RFC 814. July, 1982.

[17] Clark, David D., J. Wroclawski, K. Sollins, and R. Braden. "Tussle in Cyberspace: Defining Tomorrow's Internet." *SIGCOMM '02*, August 19-23, 2002. Pittsburgh, Pennsylvania.

[18] Clark, David, R. Braden, A. Falk, and V. Pingali. "FARA: Reorganizing the address architecture." August, 2003. ACM SIGCOMM Computer Communications Review, Proceedings of the ACM SIGCOMM workshop on Future directions in network architecture. Vol. 33 Issue 4. August, 2003. pp 313-321.

[19] Clark, Ian, O. Sandberg, B. Wiley, and T. Hong. "Freenet: A Distributed Anonymous Information Storage and Retrieval System." Accessed from http://freenetproject.org /freenet.pdf April, 2005.

[20] Comcast Terms of Service accessed at http://www.comcast.net/terms/use.jsp April, 2005.

[21] Communications Act of 1934 as amended by the Telecommunications Act of 1996. Accessed at http://www.fcc.gov/Reports/1934new.pdf April, 2005.

[22] Crowcroft, Jon, S. Hand, R. Mortier, T. Roscoe, and A. Warfield. "QoS's Downfall: At the bottom, or not at all!" *Proceedings of the ACM SIGCOMM 2003 Workshops*. August, 2003 in Karlsruhe, Germany.

[23] De Vivo, Marco, E. Carrasco, G. Isern, and G. de Vivo. "A Review of Port Scanning Techniques." ACM SIGCOMM Computer Communications Review, April, 1999, p41-48.

[24] Dolzer, Burgstahler, H. Jahnert, S. Macian, and W. Payer. "Beyond Technology: The Missing Pieces for QoS Success." *Proceedings of the ACM SIGCOMM 2003 Workshops*. August, 2003 in Karlsruhe, Germany.

[25] Egevang, K. "The IP Network Address Translator." RFC 1631. May, 1994.

[26] Extreme Networks, Inc. "Meeting the Peer-to-Peer Challenge in Higher Education – a Network Designer's View." (2003) Accessed from http://www.packeteer.com/resources/partners/EXTR_PKTR_P2P_White_Paper.pdf April, 2005.

[27] FCC Consent Decree "In the matter of Madison River Communications, LLC and affiliated companies." DA 05-543.

[28] *FCC News*. "FCC Chairman Michael K. Powell Commends Swift Action to Protect Internet Voice Services." March 3, 2005.

[29] For example, Cisco makes several routers enabled with stateful packet inspection. One example is the SOHO 90 Series Secure Broadband routers. http://www.cisco.com/en/US/products/hw/routers/ps380/products_white_paper09186a008008872f.shtml

[30] For example, see VerizonWireless calling plan options for Cambridge, MA. See http://www.verizon.com.

[31] Ford, Bryan, and P. Srisuresh, and D. Kegel. _(2005) "Peer-to-Peer Communication Across Network Address Translators." Accessed at http://www.brynosaurus.com/pub/net/p2pnat/ April, 2005.

[32] Fraleigh, Chuck, S. Moon, B. Lyles, C. Cotton, M. Khan., D. Moll, R. Rockell, T. Seely, and C. Diot. "Packet-level Traffic Measurements from the Sprint IP Backbone."

[33] Gillett, Sharon. "Universal Service: Defining the Policy Agenda in the Age of the Internet." *The Information Society*. Vol 16. No. 2 November, 1999.

[34] Grossman, D. "New Terminology and Clarifications for Diffserv.." RFC 3260. April 2002.

[35] Grossman, Wendy. "Only Connect." *The Inquirer*. December 26, 2003. Accessed at http://www.theinquirer.net/?article=13343 April, 2005.

[36] Gulbrandsen, A., P. Vixie, and L. Esibov. "A DNS RR for specifying the location of services (DNS SRV)." RFC 2782. February, 2000.

[37] Handley, M., H. Schulzrinne, E. Schooler, and J. Rosenberg. "SIP: Session Initiation Protocol." RFC 2543. March, 1999.

[38] Hernendez-Campos, F., A. Nobel, K. Jeffay, and F.D. Smith. (2003) "Statistical Clustering of Internet Communication Patterns." Available at http://www.cs.unc.edu/Research/dirt/abstracts/INTERFACE-03-abs.html.

[39] Hilden, Julia. "File Sharing goes before Supreme Court." CNN.com. February 16, 2005. Accessed http://www.cnn.com/2005/LAW/02/16/hilden.fileswap/ April, 2005.

[40] Hughs, Jeff. "Proxy Appliances Control Web Access." *Network World*. May 20, 2005. Accessed at http://www.networkworld.com/news/tech/2004/0510/techupdate.html April, 2005.

[41] Karagiannis, Thomas, A. Boido, N. Brownlee, and M. Faloutsos. "Is P2P dying or just hiding?" IEEE Communication Society, *Globecom* 2004, 2004.

[42] Karagiannis, Thomas, A. Broido, M. Faloutsos, and K. Claffy. "Transport Layer Identification of P2P Traffic." *IMC '04,* October, 2004. pp 121-134.

[43] *Kazaa.com.* “Peer-to-Peer and How Kazaa How Works.” Accessed at http://www.kazaa.com/us/help/glossary/p2p.htm April, 2005.

[44] King, Ian. "You've been served!" February 24, 2005. TerminalCity. Accessed at http://www.terminalcity.ca/index.php?option=com_content&task=view&id=260&Itemid=2 April, 2005.

[45] La Monica, Paul. “Music Industry Sues Swappers.” CNNMoney. September 8, 2003. Accessed http://money.cnn.com/2003/09/08/technology/riaa_suits/ April, 2005.

[46] Lee, Henry, and V. Thing. (2004) "Port Hopping for Resilient Networks." *Vehicular Technology Conference.* 2004 IEEE 60th. September, 2004. pp. 3291-3295.

[47] Lemley, Mark and L. Lessig. “The End of End-toEnd: Preserving the Architecture of the Internet in the Broadband Era.” 48 *UCLA Law Review*. April, 2001.

[48] Lipson, Adam. "Port blocking isn't enough for security." InternetWeek.com Accessed at http://www.internetweek.com/breakingNews/showArticle.jhtml?articleID=15500455 April, 2005.

[49] Locklear, Fred. "Comcast port 25 blocking results in less spam." Arstechnica.com. June 29, 2004. Accessed at http://arstechnica.com/news.ars/post/20040629-3945.html April, 2005.

[50] Lottor, M. “TCP Port Service Multiplexer (TCPMUX).” RFC 1078. November, 1988.

[51] Malcom, John, P. Kelley, and R. Richardson. Joint Petition for Expedited Rulemaking to Resolve Various Outstanding Issues Concerning the Implementation of the Communications Assistance for Law Enforcement Act.” Received by the Federal Communications Commission March 10, 2004. Accessed at http://www.askcalea.com/docs/20040310.calea.jper.pdf April, 2005.

[52] Merkovic, Jelena, and P. Reiher. “A Taxonomy of DDOS Attack and DDOS Defense Mechanisms.” *ACM SIGCOMM Computer Communications Review*, Vol. 34 No. 2. April 2004. pp39-54.

[53] Mockapetris, P. “Domain Names – Concepts and Facilities.” RFC 1034. November, 1987.

[54] Mockapetris, P. “Domain Names – Implementation and Specification.” RFC 1035. November, 1987.

[55] Moon, D.A.. “Chaosnet.” MIT Artificial Intelligence Laboratory Memorandum 628. June, 1981.

[56] Odlyzko, Andrew. “Internet Pricing and the History of Communications.” *AT&T Labs – Research*. Febraury 8, 2001. Accessed at http://www.dtc.umn.edu/~odlyzko/doc/history.communications1b.pdf April, 2005.

[57] Pai, Vinay and P. Rana. “A Transparent Framework for Enabling Incoming TCP Connections to Hosts Behind a NAT Gateway.” 2003. *Computer and Communications Review*. IEEE p 572-575.

[58] Peterson, Larry, and B. Davie. (2000) *Computer Networks: A Systems Approach.* Academic Press. San Diego, CA.

[59] Picture taken from http://www.enterasys.com/products/whitepapers/ssr/network-trans/netfig7.gif April, 2005.

[60] Postel, Jon. “Internet Protocol.” RFC 791. September, 1981.

[61] Postel, Jon. “Transmission Control Protocol.” RFC 793. September, 1981.

[62] Postel, Jon. “User Datagram Protocol.” RFC 768. August, 1980.

[63] PrairieWave Communications Frequently Asked Internet Related Questions. PrairieWave.com. Accessed at http://www.prairiewave.com/support/internet/portblock.htm April, 2005.

[64] Rosenberg, J., J. Weinberger, C. Huitema, and R. Mahy. "STUN: Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)." RFC 3489. March, 2003.

[65] Rosenberg, J., R. Mahy, and C. Huitema. "Traversal Using Relay NAT (TURN)." Internet Draft, work in progress,expires August 22, 2005.

[66] Saltzer, J.H., D.P. Reed, and D.D. Clark. "End-to-End Arguments in System Design." *ACM Transactions on Computer Systems,* Vol. 2, Issue 4. November 1984. pp. 277-288.

[67] Sandvine, Inc. "Peer-to-Peer File Sharing: Port Hopping and Challenges to Traffic Control Methodology." Technical Whitepaper. Accessed at http://www.sandvine.co.uk/solutions/resource_library.asp April, 2005.

[68] Sen, Subhabrata, O. Spatscheck, and D. Wang. "Accurate, Scalable In-Network Identification of P2P Traffic Using Application Signatures." *WWW 2004,* New York. Pp. 512-521. Available on ACM Digital Library, accessible at http://www.acm.org.

[69] Senie, D. "Network Address Translator (NAT)-Friendly Application Design Guidelines." RFC 3235. January, 2002.

[70] Slashdot Discussion in response to: Ulrich, Johannes, "Internet Service Providers: The Little Man's Firewall?" which has since become inaccessible online but discussed the security benefits accrued to consumers from ISP port blocking. Accessed at http://slashdot.org/article.pl?sid=03/09/07/2343254&mode=thread&tid=126&tid=172&tid=95 April, 2005.

[71] Solum, Lawrence and M. Chung. "The Layers Principle: Internet Architecture and the Law." University of San Diego School of Law, Research Paper 55. June 2003. Accessed at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=416263 April, 2005.

[72] Srisuresh, P., and K. Egevang. "Traditional IP Network Address Translator (Traditional NAT)." RFC 3022. January, 2001.

[73] *The Honeynet Project & Research Alliance.* "Know Your Enemy: Tracking Botnets." (2005) Accessed at http://www.honeynet.org/papers/bots/ April, 2005.

[74] The House Policy Committee. "Policy Statement: Establishing Global Internet Freedom." 19 September 2002. Accessed at http://policy.house.gov/html/news_item.cfm?id=112 24 March 2005.

[75] US CERT http://www.uscert.gov

[76] WhirlPool Forums, Other Broadband, Veridas, "TeeGee Port Blocking/Throttling." Accessed at http://forums.whirlpool.net.au/forum-replies.cfm?t=305067&p=1 April, 2005.

[77] Whitehouse, Steve and P. Caulfield. "DECnet Network Protocol." *Linux Journal* Vol. 1999, Issue 60es. April, 1999.

[78] WildwestDomains product offerings, linked from http://www.wildwestdomains.com/stockstores.aspx?ci=843. Product prices listed at https://www.securepaynet.net/gdshop/hosting/shared.asp?ci=1070&app_hdr=1387&prog_id=domainspricedright April, 2005.

[79] Wu, Tim. "Netowrk Neutrality, Broadband Discrimination." *Journal of Telecommunications & High Technology Law*, Vol. 2, pp. 41. 2005.

[80] Xerox Corporation. "Internet Transport Protocols." Document No. XSIS 028112. December 1981.

[81] Yoo, Christopher. "Would Mandating Broadband Network Neutrality Help or Hurt Competition? A Comment on the End-to-End Debate." *Journal of Telecommunications and High Technology Law*, Vol. 3, 2004.