(RE)DESIGNING THE INTERNET: A CRITICAL CONSTRUCTIVIST ANALYSIS OF THE NEXT GENERATION INTERNET PROTOCOL

by

Michael Felczak
B. Math., University of Waterloo, 2001

THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF ARTS

In the
School of Communication

SIMON FRASER UNIVERSITY

Summer 2005

APPROVAL

Name:	**Michael Felczak**
Degree:	**Master of Arts**
Title of Thesis:	**(Re)Designing the Internet: A Critical Constructivist Analysis of the Next Generation Internet Protocol**

Examining Committee:

Chair: **Dr. Gary McCarron**
Assistant Professor, School of Communication

Dr. Richard Smith
Senior Supervisor
Associate Professor, School of Communication

Dr. Andrew Feenberg
Supervisor
Professor, School of Communication

Dr. Robert Cameron
External Examiner
Professor, School of Computing Science

Date Defended/Approved: July 28, 2005

ABSTRACT

By means of an analysis of texts that document the design of the next generation Internet Protocol (IPng, IPv6) this thesis examines the direction and extent to which participants in this process influenced the IPng design. The analysis suggests that many IPng design participants were commercially oriented. The military and groups that focused on IPng support for particular technologies also took part. In contrast, state and civil society participation was notably non-existent. The problems and concerns of commercial, military, and technology oriented groups were largely accommodated by the eventual IPng design. Issues of focus traditionally foregrounded by the state and civil society did not inform the design process. Initial IPv6 implementations will clarify and define in greater detail the IPv6 design. Given the potential of IPv6 to drastically alter the Internet and its use, this implementation process presents a strategic opportunity to define the direction of Internet development.

DEDICATION

To my grandfather, Adam, and my great aunt Stefania, both of whom fought with serious illness during the writing of this thesis and who define strength of mind, determination, and perseverance.

ACKNOWLEDGEMENTS

The completion of this thesis would not have been possible without the help and generosity of many people. First and foremost, I would like to thank Richard Smith for his guidance, constructive feedback, and warm support throughout the duration of this degree. Richard's encouragement, suggestions, and open-mindedness provided an ideal environment for the exploration of ideas and the writing of this thesis.

I am also greatly indebted to Andrew Feenberg, whose commitment and interest in this project contributed to an intellectually rewarding experience. I am especially grateful for the many discussions and conversations that helped me to develop and improve the arguments in this thesis.

I would also like to thank everyone who helped me with earlier drafts of this thesis, which was greatly improved as a result of many stimulating conversations and the editing and proofreading efforts of many people. I would like to especially thank my brother, Luke, for his encyclopaedic knowledge of most things, his willingness to explore ideas, and his patience during the final stages of editing and proofreading. Likewise, I would like to thank Ted Hamilton for his pointed suggestions and help with the final draft given a tight schedule.

Last but certainly not least, I would like to thank my family for their endless support and encouragement. The completion of this thesis would not have been possible without the unconditional support of my parents and their understanding and accommodation of the needs inherent to this activity. I would also like to extend a special thank you to my grandmother, whose good humour and rich cooking helped me immeasurably during an emotionally challenging time.

TABLE OF CONTENTS

Approval ii

Abstract iii

Dedication iv

Acknowledgements v

Table of Contents vi

List of Figures viii

List of Tables ix

List of Abbreviations x

Glossary xi

Chapter 1: Motivation, Context, and Relevance 1

- Introduction 1
- Internet Change and Development: Two Examples 3
- Conceptualizing the Internet: Collections, Layers, and Protocols 6

Chapter 2: Ontology and Epistemology 10

- Introduction 10
- A Sociology of Knowledge and Technology 11
- Humans, Non-Humans, and Heterogeneous Engineering 13
- Implications, Processes, and Irrelevant Social Groups 17
- Critical Constructivism 20
- Hegemony, Technical Codes, and Technical Change 21
- Agency, Locality, and Representation 23
- Research Questions 25

Chapter 3: Methodology 27

- Introduction 27
- Actors, Networks, and Actor-Networks 28
- Text, Description, and Explanation 29
- Internet Standardization 30
- Selection Criteria and Procedure 36
- Interpretative Framework 38
- Validity, Reliability, and Generalizability 40

Chapter 4: Findings 43

- Introduction 43
- The IPng Area and RFC 1550: Interesting and Enrolling Others 44
- The Interested Parties 45

The IPng Area and RFC 1726: Translating Everyone 83
Expectations and Specifications 89
Accommodating Programs and Anti-Programs 94
Conclusion 97

Chapter 5: Discussion 99

Introduction 99
Irrelevant Social Groups 100
Interoperability, Open Standards, and the Future of New Media 104
IPv4, Inertia, and Expectations 105
Formal Bias 109
Technical Codes, Implementation, and Sociotechnical Change 111
Regrets and Reservations 115
Future Research 116
Conclusion 117

Reference List 119

LIST OF FIGURES

Figure 1: The Internet According to Benkler's Model of a Communication System..........7

LIST OF TABLES

Table 1: Addressing 92
Table 2: Routing 92
Table 3: Transition 92
Table 4: Security 93
Table 5: Other Protocols 93
Table 6: Link and Network 93
Table 7: Packets 93
Table 8: Miscellaneous 94

LIST OF ABBREVIATIONS

ARPANET	Advanced Research Project Agency Network
ATM	Asynchronous Transfer Mode
BOF	Birds of a Feather
CLNP	Connectionless Network Protocol
DNS	Domain Name System
IAB	Internet Architecture Board
ICMP	Internet Control Message Protocol
IESG	Internet Engineering Steering Group
IETF	Internet Engineering Task Force
IPng	Internet Protocol Next Generation
IPv4	Internet Protocol, Version 4
IPv6	Internet Protocol, Version 6
IPSEC	Internet Protocol Security
OSI	Open Systems Interconnection
P2P	Peer-to-Peer
PPP	Point-to-Point Protocol
RFC	Request for Comments
RSVP	Resource Reservation Protocol
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
WWW	World Wide Web

GLOSSARY

Datagram A self-contained data packet that contains addressing information.

Domain Name System An Internet service and mechanism that enables the translation of a host name into an IP address and vice-versa.

Globally-unique address An address that uniquely identifies a node across some set of interconnected networks.

Hierarchical address An address that not only identifies a node in a network, but also includes information regarding its location in a hierarchical network. For example, an IP address includes information regarding the network and subnet to which a host belongs to.

Hierarchical routing Routing that has been subdivided into smaller networks and structured as a hierarchy. Each level in the hierarchy is responsible for its own routing.

Latency The amount of time that it takes a packet of data to reach a network destination point.

Locally-unique address An address that uniquely identifies a node with respect to a particular (local) network.

Multicasting The ability to send information to multiple receivers simultaneously.

Multi-homing A single network interface with multiple addresses. Alternately, a node with multiple network interfaces.

Network topology The physical and logical arrangement of nodes in a network.

Policy routing Routing that is based on some user-defined policy that specifies particular routes for particular applications, data traffic, or users, for example.

Reliability The ability to maintain some predetermined level of network performance given expected network conditions.

Resource reservation The ability to reserve network resources for the transmission of data.

Route aggregation The identification of multiple routes through a single router entry. Route aggregation takes advantage of a common address prefixes in hierarchical addresses. Rather than keeping track of individual entries for all nodes on another network, a router can simply maintain information on how to reach the network, which is then responsible for further routing.

Scalable addressing An addressing system that is able to accommodate all existing network nodes as well as new nodes that may be added in the future.

Scalable routing A routing system that is able to accommodate all existing network nodes as well as new nodes that may be added in the future.

Service selection The ability of a network to deliver multiple services. For example, a network may provide best-effort service, with no delivery guarantees, as well as some form of differentiated service that may provide guarantees with respect to delay, error rates, or other characteristics.

Source routing Sender indicates the path through the network by identifying nodes between itself and the receiver.

Topology independent address allocation The allocation of network addresses based on considerations other than network topology. For example, address allocation based on geography.

CHAPTER 1:
MOTIVATION, CONTEXT, AND RELEVANCE

Introduction

Since its initial public appearance the Internet has been the focus of much thought and speculation both inside and outside of academic circles. The potentials and infinite possibilities so often cited leading up to the dot-com boom have faded since the dot-com bust and have been replaced by more sober predictions and forecasts for this still new medium.

Today, despite continued persistence and prediction of unparalleled changes in some circles, many commentators and the public are more sober about the Internet. An infrastructure once developed and run almost exclusively by non-commercial groups and organizations has increasingly become commercialized. The web, once a cyber-place for the technically curious is now occupied by the largest multinational corporations in the world. E-commerce is expanding at a rapid rate and appears to be here to stay. Despite these developments new Internet applications, both commercial and non-commercial, continue to be developed. For example, peer-to-peer content sharing and voice-over-IP, or Internet telephony, have become increasingly popular and have attracted the (often negative) attention of the entertainment and telecommunication industries.

The conflicts and controversies that have arisen around new media such as peer-to-peer networking and Internet telephony not only foreground the political and economic relations that are threatened by these media, but also make explicit taken for granted conceptions necessary to preserve these relations. Notions regarding intellectual property, copyright, and the right to communicate have become the subject of recent public and policy debate.

The legal and policy outcomes of these debates will to some extent define the boundaries for new media. From a technical standpoint, however, Internet developers continue to develop new media that utilize the Internet's existing infrastructure. Given this activity, it seems difficult to reach agreement and closure regarding the Internet's future. One can certainly analyse existing Internet applications, such as the web, email, or instant messaging, but the generalization from these specific media to the Internet as a whole is not as easy as it seems given that new media forms continue to be developed by various groups and organizations that span a spectrum from non-profit open source advocates to established entertainment conglomerates.

The Internet architecture, which makes this dynamic possible, has been relatively unchanged since its early inception in the Advanced Research Project Agency Network (ARPANET). At the time, the academic and military community involved in its development selected a simple design that was able to efficiently utilize the existing communication and computing infrastructure and withstand network outages and disruptions. In brief, the architecture works as follows: a message is broken up into smaller components, known as packets, which are transmitted via various network paths to their destination and re-assembled by the receiver to obtain the original message. Simply stated, the architecture provides a simple data transport service to users and their applications.

Although this basic architecture has been the foundation of the Internet for over three decades, in recent years there has been a concerted effort to redesign this architecture. Groups and organizations for which the architecture is problematic have voiced their concerns and argued that changes need to be made to better accommodate their needs and conceptions of the Internet. Some of these visions include an Internet where users are charged for the individual packets that they are responsible for originating. Likewise, certain conceptions require that the Internet provide different grades of service for different grades of users. Those who are willing to pay premium prices should be guaranteed premium network services, even if this implies that nonpremium Internet traffic must be delayed or undelivered. These conceptions clearly understand the Internet, communication, and users in a very particular way such that the existing Internet architecture is ill-suited to these conceptions.

The purpose of this thesis is to examine in detail the various groups and organizations that participated in the redesign of the Internet's architecture, to consider how effective these groups were in translating their demands into specific design features of this architecture, and to reflect on the implications of these features for the next generation Internet and new media. As I argue in the remainder of this chapter, this area of research has been largely ignored by Internet researchers in particular and communication scholars more generally. An understanding of the issues and politics that underlie the Internet's architecture is central to critical engagement and strategic intervention into these issues.

Internet Change and Development: Two Examples

To foreground in greater detail the relation between the Internet's architecture and new media, I begin with an analysis of two samples of research that attempt to come to terms with recent developments on the Internet, but that ultimately fail to adequately theorize the implications of these developments as a result of a poor understanding of this relation. Specifically, the research of Korinna Patelis and Dan Schiller, while contributing to our understanding of recent Internet commercialization, illustrate some common assumptions and arguments that appear in analyses of the Internet.

To start, consider the analysis of Patelis (2000), who challenges the hype and ideology of the Internet, which she terms "Internetphilia," and which she believes is rampant throughout our society. According to proponents of Internetphilia, the Internet will usher in a new era whose features are fundamentally different from those of the past. The Internet represents a clear break with the past, since its virtual embodiment is free from the reality of social, economic, and political processes.

Suffice it to say, Patelis provides a much needed critique of Internetphilia. She argues that economic, political, and social realities have manifested themselves in cyberspace, which is not immune from their influence. Patelis analyzes geographic patterns of Internet host concentration, Internet access inequalities, Internet backbone dispersion and connection quality as well as usage and development patterns that favour corporate users over public users (pp. 91-97).

Despite this contribution, when Patelis's analysis is examined in terms of its conceptual understanding of the Internet, it becomes apparent that Patelis mixes and interchanges the terms "Internet" and "web" and ultimately reduces the former to the latter. As an example, consider the following:

> The Internet is a commodified medium: the exchange value of on-line communication is prioritised over its other values. There are three main ways in which such prioritisation takes place, each generating significant revenues. First, the commodification of access ... Second, the commodification of Internet navigation tools and search engines ... Third, advertising: Web advertising revenue grew by 28 per cent to US $169.8 million ... (p. 91)

The passage begins with reference to the commodification of the Internet and the prioritization of exchange value for online communication. The first example, Internet access, is consistent with this claim, but the second and third points focus almost exclusively on the web. Both search engines as well as web advertising apply only to the web and not necessarily to the Internet as a whole. That is, Patelis analyzes recent web developments, provides evidence and examples of commodification on the web, and then generalizes these specific instances to the Internet as a whole.

A similar logic underlies Patelis's analysis of e-commerce (pp. 96-97). Patelis cites the recent emergence of cybermalls and e-commerce web sites and uses this as evidence to conclude that the Internet as a whole is one large supermarket. For Patelis, e-commerce "shifts the function of the Internet from a commodified medium for communication to a supermarket" (p. 96). The general logic is the same: developments specific to the web are generalized to the Internet as a whole.

The tendency to reduce the Internet to the web and to posit a single, concrete, and decided function for the Internet is by no means isolated to Patelis's analysis. A somewhat similar pattern underlies Dan Schiller's (1999) account of Internet development. Tracing the policy and technology changes that long preceded the emergence of the Internet, Schiller argues that neoliberal policy has largely succeeded in shaping network technology to serve its needs. He roots the seeds of the movement in corporate U.S. pressures to liberalize data networking and telecommunication and shows how a build-up of inertia transgressed American borders to influence policies and developments abroad.

Although Schiller's analysis is particularly strong in its coverage of policy changes, corporate network development and mobilization, and telecommunication system restructuring, Schiller nonetheless suggests that the Internet, its function, and its fate is largely sealed and determined by the corporate and military powers that gave it birth. Like Patelis, Schiller's analysis of recent Internet developments is also largely confined to the web, which is analyzed in the context of online education and e-commerce. Although the suggestion is more implicit than that in Patelis's analysis, Schiller nonetheless suggests that recent commercial developments in these two online domains may be generalized to the Internet as a whole, which is posited to exclusively support the U.S. neoliberal project.

The reduction of the Internet to the web is not merely a problem for engineers and computer scientists to fret and squabble about. The difference between the two entities has important implications for examinations and analyses that try to come to grips with the Internet and recent developments. The Internet and the web are by no means reducible to each other, despite the current popularity of the web.

Conceptualizing the Internet: Collections, Layers, and Protocols

Patelis's and Schiller's analyses, while acknowledging the many components that make up the Internet, privilege particular components over others and use these components to generalize the Internet as a whole. That is, both view the Internet as a collection of networks and the hardware, routers, backbones, cables, software, protocols, and applications that make it work. From this collection, certain components are privileged, analyzed, and then used as support to make general claims about the Internet. Conceptualizations that fail to explore the Internet and the many interactions between its components work with what could be called a "collection" model of the Internet. That is, the Internet is conceptualized as a broad collection of parts that include hardware, software and applications, but the various relations and dependencies of the component parts are not adequately accounted for.

In light of this observation, it seems worthwhile to accept the evidence provided by Patelis, Schiller and others who correctly point out that many changes have taken place in the Internet's infrastructure and usage. At the same time, we can question some of the conclusions that these analyses point to by critically examining the authors' conceptualizations of the Internet and their failure to consider thoroughly the various areas of Internet development.

It is the contention of this thesis that the protocols at the heart of the Internet's architecture and the social groups that define them have been largely ignored by Internet researchers. These protocols have not been adequately considered in analyses of the Internet and this omission has in part contributed to often contradictory predictions regarding future trends and patterns. An analysis of the Internet needs to address these shortcomings and clearly articulate the importance of these protocols and their relation to Internet development and change.

A good starting point for this task is Yochai Benkler's "layer" model of a communication system, which Lawrence Lessig (2002) utilizes in his analysis of the Internet. As Lessig points out, Benkler's model "helps organize our thought about how any communications system functions. But in organizing our thought, his work helps show something we might otherwise miss" (p. 23). In our case, Benkler's layer model makes explicit the relations between various Internet components and highlights the importance of protocols in the Internet's design and function.

Following the technique of network designers, Benkler suggests that we view a communication system such as the Internet as divided into three distinct layers. The first layer, which is at the bottom, is the "physical" layer and in the case of the Internet it contains the cables, wires, and the physical infrastructure that links computers together. The middle, or "logical" layer, includes the code and logic that makes the hardware run. It includes protocols and specifications and the software that implements them. Lastly, the top, "content" layer includes the content that is transmitted from computer to computer and the user applications that enable the transmission of this content. Figure 1 illustrates Benkler's model in the case of the Internet.

Figure 1: The Internet According to Benkler's Model of a Communication System

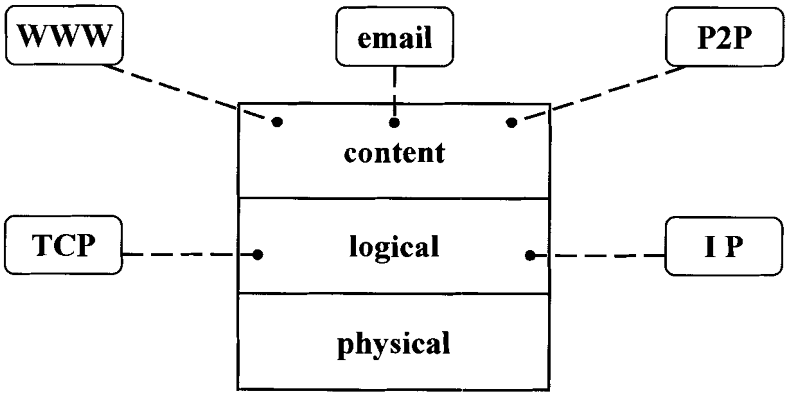

What is evident in Benkler's model is the hierarchical relation between the various components. As Figure 1 illustrates, the content layer is dependent on the logical layer, which, in turn, is dependent on the physical layer. Benkler's model applied to the Internet highlights the central place of the logical layer and the data transport and transmission protocols that it contains. Both the Transmission Control Protocol (TCP) as well as the Internet Protocol (IP) reside in this layer. In order to function, the logical layer needs to be able to interface with the physical layer. Likewise, applications at the content layer need to understand and follow the conventions of the logical layer in order to transmit content across the Internet.

Returning to our previous example, changes to the web, which resides in the content layer, whether they involve increased commercialization or alternative developments, can take place and the web will continue to function as the web as long as the specifications that define it continue to satisfy the requirements of the logical layer. A similar argument can also be made for email or peer-to-peer networking technologies, for example. Likewise, as long as changes at the physical layer are compatible with the basic assumptions with which the protocols at the logical layer work with then the logical layer and the applications that it supports can continue to function.

Stabilization of technologies can occur in multiple places. First, the applications at the content layer are designed and developed as separate from each other. That is, email, the web, and instant messaging, for example, are unique technologies that have been developed at various times by various individuals, groups, and organizations. In principle, the applications are isolated and changes to the web do not in and of themselves imply changes to email or instant messaging.

Within the logical layer, the TCP and IP protocols define the mechanisms that enable applications at the content level to transmit data across the Internet. This technology is as rich in meaning and significance as the popular applications at the content layer. However, since it appears as void of meaning and norms to most people, it is often overlooked as a site of contestation. Particular normative considerations embodied in the design of these protocols have had important consequences for the nature and quality of the Internet that we know today.

As this thesis will illustrate, there are many ways in which this transport of data can take place. Should users need to provide proof of their identity before they are able to use the network? Should the network treat all applications equally or should certain applications, such as pay-per-view information and entertainment, receive priority over other, less profitable forms of communication? Should network service providers charge users by the packet instead of by the hour? These are pressing questions. The design of the next generation Internet Protocol has been taking place since the early 1990s and various proponents have replied affirmatively to these propositions. This design process, its politics, and its outcome, a protocol specification that defines the boundaries of possibility, is the focus of this thesis.

CHAPTER 2:
ONTOLOGY AND EPISTEMOLOGY

Introduction

In this chapter I outline and successively refine the ontological and epistemological position of this thesis. I begin with an analysis of social constructivist approaches to technology study as exemplified by Pinch and Bijker (1987) and while I retain some of the key insights and conceptualizations of this position I reject the exclusive privilege granted to the social entailed by these approaches. Instead, I consider the less restrictive position offered by actor network approaches that treat symmetrically social and material factors in the success and failure of technology. Specifically, I examine Latour's (1987) contribution to science and technology studies and outline the implications entailed by this approach to the study of technology and to our understanding of the relation between technology and society. I then shift focus to consider some of the criticisms levelled at actor network and social constructivist approaches. Langdon Winner (1993) and Andrew Feenberg (1999) in particular have been vocal regarding the shortcomings and inadequacies of these formulations. I analyze the criticisms and issues of contention and move to outline in detail Feenberg's critical constructivist approach, which addresses some of the shortcomings of actor network and social constructivist approaches. By acknowledging the political dimension of technology and the role that it plays in contemporary societies, Feenberg locates technology at the intersection of social, political, and economic circumstances and draws out the implications entailed by this locus of concentration. Lastly, in order to interweave the ontological and epistemological commitments entailed by the above analysis I conclude with a detailed formulation of the research questions that guide this thesis.

A Sociology of Knowledge and Technology

Trevor Pinch and Wiebe Bijker (1987) integrate insights from the sociology of scientific knowledge and social construction of technology to develop a conceptual and methodological framework for the study of technology and its design. According to Pinch and Bijker, the study of science and the study of technology have many important similarities and overlaps that are often unacknowledged by viewpoints that consider each a distinct and unique sphere of activity (p. 17).

Pinch and Bijker's social construction of technology examines the development process of technology as an alternation of variation and selection. This multidirectional view highlights the various and often differing development paths of technology and considers the selection mechanisms that give rise to the eventual stabilization of technological designs. In contrast to historical hindsight examinations, which support the illusion of a single, linear path of development, multidirectional models seek to illustrate the non-linear and sometimes conflicting designs and patterns of development (p. 28).

The multidirectional model accounts for successful as well as ultimately unsuccessful designs. In order to map and understand these successes and failures, relevant social groups are identified, since social groups involved with a technology give it meaning and constitute problems for which the technology may provide potential solutions. As Pinch and Bijker point out, "a problem is defined as such only when there is a social group for which it constitutes a 'problem'" (p. 30). Relevant social groups may include institutions and organizations or unorganized groups of individuals. The key requirement is that all members of the group share the same set of meanings with respect to a particular technology.

Borrowing the term from the sociology of knowledge, Pinch and Bijker define the "interpretative flexibility" of a technological artefact as the variable and socially constructed interpretation of a technology and its design (p. 40). That is, there is not just one possible way of designing a technology. The meanings and understandings of relevant social groups define both what the technology is for the groups as well as how it should be best designed.

An awareness of the problems each group has with respect to the technology is complimented by an identification of potential solutions available by means of new or modified designs as well as through alterations of legal structures, regulations, and social norms (pp. 35-39). An analysis of the various problem/solution claims may lead to the identification of a wide variety of conflicts between the relevant social groups. These conflicts may include conflicts over technical requirements that address certain concerns at the detriment of others, conflicting solutions to the same problem, as well as social or moral conflicts that arise from particular designs (p. 35).

The multidirectional model elaborated by Pinch and Bijker draws attention to the growing and diminishing degree of stabilization of competing designs. This degree of stabilization is, in principle, different for different social groups. In the case of the bicycle, Pinch and Bijker's example, it is clear that invention and innovation are not isolated events. It took 19 years before the features and characteristics of the "safe bicycle" were deemed as essential elements of its design. Prior to this stabilization, a wide range of bicycles and tricycles of varying designs competed as potential solutions to the safety problem, which was expressed by older individuals, women, and riders who sought safe transportation in favour of the speed and sport provided by "high-wheeler" designs (p. 39).

For Pinch and Bijker, stabilization of an artefact involves the "disappearance" of problems and the closure of a technological "controversy." That is, closure need not require that a problem actually be solved in the common sense understanding of this term. The key point is that the relevant social groups "see the problem as being solved" (p. 44).

In sum, Pinch and Bijker's approach foregrounds the interpretative flexibility of technologies and the controversies that arise as various social groups attempt to shape and direct multiple paths of development. Different interpretations lead to different problem/solution claims and to varying degrees of stabilization over time. Closure mechanisms give rise to the disappearance of problems and signal the end of technological controversies.

Humans, Non-Humans, and Heterogeneous Engineering

The social constructivist approach highlights the underdetermined nature of technology and argues that technical or natural factors are insufficient to explain the shape and form of technologies. The principle of symmetry from the sociology of science is extended to the sociology of technology: the same type of social explanation needs to be used for both successful and unsuccessful technologies. The tendency of explaining successful technologies by referring to their inherent characteristics and explaining unsuccessful technologies by alluding to social factors and social conditions needs to be abandoned.

Although extension of the principle of symmetry from science to technology ensures consistency of explanation across technology successes and failures, it nonetheless privileges the social above all other factors in the explanation. Extending the principle of symmetry further, actor network approaches adopt what is often called the "generalized principle of symmetry," which states that the same type of explanation should be used for all of the elements that make up a technology, regardless of whether these elements are social, natural, or technical. That is, the same type of explanation should be used for both successful as well as unsuccessful technologies, but the social elements should not be privileged and given special status over other elements.

As John Law (1987, p. 113) argues, it is not that the social is not important, since it often is. Social factors are at times the dominant factors in the growth and stabilization of a technology. Instead, the issue is not to privilege the social at the outset, since at times technical or other elements may resist all efforts to shape and control them, and as a consequence, these elements may better explain the success and shape of the resulting technology.

As an example that illustrates the role of various human and non-human elements, Law (p. 112) refers to Thomas Edison and the challenges that Edison faced when he attempted to develop the light bulb. Edison needed to consider how to supply electric lighting at an affordable price (economic), how to persuade policy makers to permit the development of a power system (political), how to achieve the best balance between the length of power lines, current flow, and the amount of voltage (technical), and how to find a high resistance incandescent filament for his bulb (scientific). In a case such as this, it is difficult to privilege exclusively the social, which is closely tied and interrelated with the technical and the scientific. This is not to say that science is not social, but simply to point out that filaments may or may not provide Edison with the high resistance that he requires. Proponents of actor network approaches argue that analyses must account for a multiplicity and variety of factors and elements, some of which may be non-human.

According to Latour (1987), a technology has elements tied to it that make it more or less "real." If Edison can successfully master the high resistance incandescent bulb, direct policy makers, and conceive of an efficient and affordable power grid then electric lighting becomes more and more "real." On the other hand, if policy makers are opposed or if experiments into high resistance bulbs do not yield promising results or if electricity cannot be delivered affordably into the home then electric lighting becomes less "real," since it can only exist as blueprints and mathematical formulas on Edison's desk. For a technology to become "real," the elements that make it up need to be successfully juxtaposed and controlled and the result needs to be turned into a "black box" that conceals its complexity and is incorporated by others as an unproblematic, unified whole.

According to Latour, builders of technology need to do two things in order to be successful. First, they need to enrol others so that they participate in the construction. And second, they need to control the behaviour of those enrolled so that their actions are predictable. That is, for Latour there is a set of strategies to enrol and interest human actors and a second set of strategies to enlist non-human actors so as to control and hold the human actors in place (p. 132).

As with social constructivist approaches, the notion of interest is central to an understanding of these mechanisms. However, unlike social constructivist approaches, Latour acknowledges the explicit interests of social groups and individuals and suggests that these explicit interests may be done away with or modified in the process of technology construction (p. 114). The notion of "translation" is central to this conception. If we understand technology builders as "heterogeneous engineers" (Law, 1987) that try to interrelate and associate human and non-human elements into durable relations, then "translation" may be understood as the interpretation given by heterogeneous engineers of their own interests and that of the actors that they try to enrol to further their efforts (Latour, 1987, p. 108).

This is not to say that explicit interests do not need to be considered or that they do not play an important role. For Latour, technology builders utilize a variety of strategies that work with explicit as well as displaced or newly conceived interests. Translations enable technology makers to relate and link their interests with those of others, who are needed to take up and incorporate the technology as a unified and unproblematic whole.

The process of interesting and involving others is necessary but insufficient, since enrolled groups are difficult to control and may always lose interest or behave in unpredictable or undesirable ways from the viewpoint of the technology builder. To prevent these possibilities, heterogeneous engineers need to tie the invested groups to as many elements as is necessary to resist any and all efforts to disassociate the assembly. That is, the association of human and non-human elements needs to be strong enough to resist all "trials of strength" to which it may be subjected (p. 122).

For Latour, this association of elements, or actor-network, is only as strong as its weakest link, regardless of how strong some of its elements may be. Thus, if trials of strength break apart an existing assembly, network builders need to “shift the system of alliances” to compensate for the weakness and reorganize the actor-network: new elements may need to be brought in; existing elements may need to be given up; associations may need to be reorganized. Choices have to be made in such a way that the resulting actor-network resists all trials to break it apart.

Technology makers will be most successful when they are able to tie enrolled groups to non-human elements by means of various strategies of translation. The mobilisation of non-human elements, whether they are “facts,” atoms or catalysts in fuel cells, is a powerful and effective strategy, since non-human elements have been made durable and stable through past controversies and are difficult to challenge (p. 128). By tying interested groups to durable non-human elements, technology builders make it difficult for the groups to disband.

Having examined some of the strategies available to technology builders to interest others and to control their actions, we can now consider the construction of technology over time. Latour suggests that we can observe this process from two angles, or by considering technology in terms of its “sociogram” or its “technogram” (p. 138). In the former, one looks to see who the technology is designed to enrol. In the latter, one examines what the technology is tied to so as to make this enrollment inescapable. The notions of sociogram and technogram foreground the dynamic nature of technology as it changes over the course of its life.

For example, in the course of its construction, the Diesel engine changed from a blueprint accompanied by some patents, to a prototype that did not work, to a few modified prototypes that were eventually made to work, to finally, and unproblematic black box that could be sold and used in cars and trucks. In the course of its life, the Diesel engine needed to be taken up by engineers, technicians, managers, salespeople, and eventually users (p. 138). Thus, what the technology is composed of depends on who it needs to convince so that they take it up and further its development. Latour states the relationship as follows: “Each modification in one system of alliances is visible in the other. Each alternation in the technogram is made to overcome a limitation in the sociogram, or vice versa" (p. 139). Once Diesel has interested a car manufacturer to build a prototype, he needs to associate his blueprints and patents with the engineers to ensure that they build the prototype according to his design. Similarly, once a working prototype is available and the engineers are firmly in place, a new system of solid injection will need to be developed in order to turn the engine into a black box that may then be bought by consumers. Understanding what the technology is and who the people are, that is, the technology's technogram and sociogram, are two sides of the same coin. Information on one system of alliances is information on the other, and vice versa.

Implications, Processes, and Irrelevant Social Groups

In his review of the sociology of technology, Langdon Winner (1993) argues that constructivists, including proponents of actor network approaches, fail to address key issues and questions that have been the focus of much philosophy of technology and social science. In the following, I consider the merits and contributions of Winner's critique and argue that his merging of social constructivist and actor network approaches blurs important differences between the two formulations such that while some criticisms may be applicable to one approach they do not seem to apply to the other. I finish by suggesting that Andrew Feenberg's (1999) critical theory of technology addresses the shortcomings of both approaches and that it is well suited to a politically informed study of technology.

To begin with, Winner takes issue with constructivists' simplistic approaches to and considerations of relevant social groups. For Winner, the constructivists' conception of social process resembles the essential conceptions of theories of political pluralism, which foreground the interactions of various interest groups in society and their influence in decision and policy-making processes. Critics of pluralist theories argue that it is important to take note of not only the groups that are present, but also the groups that are absent and, consequently, the issues that are never articulated (Winner, 1993, p. 369). Social, political, and economic structures enable and constrain social groups and their ability to articulate as well as legitimate their concerns and issues. An analysis that considers the social process of technological design needs to consider equally which groups and issues are present as well as which groups are not present and which issues never qualify for the official agenda.

Secondly, Winner argues that constructivism fails to entertain the possibility that the social process of design may include dynamics beyond those revealed in the immediate scope of problem definition and solution formulation. As Winner points out, "[i]nsofar as there exist deeper cultural, intellectual, or economic origins of social choices about technology or deeper issues surrounding these choices, the social constructivists choose not to reveal them" (p. 371). That is, constructivism focuses on those aspects of social process that involve the construction of meanings and technologies, but fails to consider the possibility that this activity may be animated by other processes in society. For Winner, wider social, political, and economic dynamics are hidden by constructivist accounts, since the focus is exclusively directed at social actors, their interpretations, and their interactions.

Lastly, Winner is particularly critical of constructivists' general lack of moral or political commitment in their study of technology. Although he acknowledges that limiting one's scope of analysis to a descriptive exploration of technology is a worthwhile pursuit, he argues that it is necessary to eventually move beyond the merely descriptive and evaluate technological choices and decisions according to some value-based criteria. In his view, constructivists have chosen to remain silent in this regard.

Although Winner's criticisms are valuable, they unnecessarily merge social constructivist and actor network approaches and in so doing conceal key differences in conception of social process offered by these respective formulations. That is, while the actor network approach adopted for this thesis acknowledges the interpretative flexibility of technology and the underdetermined nature of the development process, it goes beyond meaning construction and the understandings of relevant social groups to consider social and material arrangements and the attempts of technology builders to interweave diverse elements in stable and durable ways. It is only through this understanding of elements and their associations that it is possible to evaluate the actor-network and identify weak elements and weak associations. This identification is a necessary prerequisite for any critical strategy of intervention aimed at exploiting these weaknesses and re-arranging the actor-network.

What is particularly problematic, however, is that actor network approaches insist on a strict methodological frame that prohibits the introduction of Winner's irrelevant social groups by the researcher. Only actors, understood in terms of activity, action and relevance with respect to the actor-network, should be accounted for and included in the analysis. As far as the actor-network is concerned, irrelevant social groups simply do not exist, since they do not influence or in any way make their presence felt on the network. This commitment undermines the critical basis of actor network approaches and needs to be acknowledged and reconciled in order to avoid accounts of technological development that normalize and, thus, support prevailing forms of domination (Feenberg, 2002, p. 31). A critical approach must account for and consider both the groups revealed by the actor-network as well as Winner's irrelevant social groups. Although this may at first seem like an incompatible requirement, I would suggest that it may be reconciled by a reflexive researcher who provides an evaluation of the actor-network and the properties, distributions and relations of its actors.

It should also be noted that actor network approaches acknowledge the explicit interests of social groups and move beyond a limited consideration of the interpretation of particular technologies. In fact, Latour's (1987) formulation includes a rich and complex treatment of these interests and the ways in which they are mobilised, transformed, and displaced during the development process. The approach acknowledges that efforts may be animated by economic, political, and social motivations but simultaneously maintains that these motivations alone cannot be relied upon to explain why some technologies succeed and others fail. What is needed is a consideration of the ways in which technology builders attempt to associate and hold together heterogeneous elements and the ways in which these assemblies either stay firmly in place or disassociate and drift apart.

Although Winner's criticisms do not apply equally to social constructivist and actor network approaches, Winner's concerns regarding the implications of technological choices do address a neglected aspect of both formulations, which tend to focus on short-term processes at the detriment of considering the longer-term commitments of technology. That is, constructivists fail to consider the implications that certain technological designs have on people's quality and experience of life and the commitments that accompany certain technological choices in terms of power, control, and ecological and social well being. By focusing exclusively on the process of design, constructivists ignore technology's contribution to defining the contours and features of the environment in which it is mobilised. Winner's critique foregrounds the need to consider technology's contribution to the definition of people's experiences and quality of life in terms of a multiplicity of commitments ranging from autonomy and control, environmental requirements, and sociotechnical relations.

Critical Constructivism

Like Winner, Andrew Feenberg (1999) acknowledges the contribution of constructivist approaches to our understanding of technology but remains skeptical with regards to constructivists' exclusion of the political dimension of technologies and their designs. To remedy this tendency, Feenberg develops a philosophy of technology that foregrounds the political character of technology, its design, and its use. That is, to the extent that needs and interests are socially defined, competing solutions are realized in different technical choices that may have drastically different requirements, potentials, and side effects (p. 84). As with other political domains, what is at stake is the conceptualization and realization of competing visions of society and the relations entailed by these realizations. Understood in this way, technology may either enforce and support the power divisions and social relations of advanced capitalist societies or undermine the social hierarchy through sociotechnical alliances that strengthen subordinate positions.

Although Feenberg is sympathetic to concerns regarding the tendencies of contemporary technological systems, which minimize meaningful human activity and reduce the environment to raw materials in their implementation of goals and rational efficiency, concerns raised in various forms by Heidegger, Ellul, and Habermas, Feenberg distinguishes these tendencies as a historical product that lacks a universal basis. In so doing, Feenberg's philosophy of technology foregrounds the political nature of technology and its design, which is a site of struggle for social groups that define what the technology is and what development it should take.

Feenberg's philosophy of technology is an attempt to account for and elaborate on the democratic interventions of experts and the lay public into technological affairs. Although at the current time this process is intermittent and includes a wide variety of tactics and resistances, these struggles draw attention to the political nature of technology and undermine technocratic justifications for the existing technological order. For Feenberg, the progress of democratic society is intimately intertwined with technological progress, which only occurs when the concerns and values of those enrolled in society's many and pervasive technological networks are embodied in the designs that structure and mediate their everyday lives.

Hegemony, Technical Codes, and Technical Change

Consistent with constructivist approaches, Feenberg extends semiotics to the study of things, but distinguishes this extension across two dimensions (pp. 84-87). The first dimension corresponds to the attribution of specific values, properties and characteristics that define what a technology is and how it should function. The second dimension corresponds to what Feenberg calls the "cultural horizon" of technology, or the cultural assumptions and common sense that forms the unquestioned background of everyday life.

If we understand hegemony as domination that seems natural to those who are dominated and which is rooted in the unquestioned background of everyday life, then it is necessary to consider shared cultural norms and values and their contribution to designs that lend material support to the prevailing hegemony. The role of a critical theory of technology is to uncover this cultural horizon and foreground contingency in the seemingly necessary (p. 87).

The relation between the cultural horizon and technological design is not unidirectional. Technologies, or more specifically, technological "regimes," contribute to the cultural horizon by providing standard ways of conceiving of problems and solutions and of legitimating certain norms and practices. That is, over time successful designs and technologies provide the basis for further development by leaving an impress on scientific knowledge and institutions, engineering practices and procedures, and standardized definitions of technologies. Technological regimes conceal the social and cultural contingencies of technologies in a neutral and highly technical language that appears self-evident to all those divorced from the specific historical conditions responsible for its particular content and form.

Feenberg introduces the notion of "technical code" to capture this relation between seemingly neutral and purely technical aspects of technological regimes and the particular social values and conditions that gave them birth (p. 88). Technology appears apolitical as long as it is conceived in terms of its technological regime, which internalises past conflicts and decisions in purely technical language and practice. The notion of technical code is an attempt to capture this movement from controversy, conflict, and historically specific circumstances to self-evident, obvious, and seemingly necessary aspects of technological regimes.

Understood in terms of technical codes, technical change is constrained and enabled by past decisions and commitments. To the extent that technical codes specify the very definition of a technology, controversy arises when attempts are made to alter this definition. To preserve and protect existing designs, technologists may mobilize technical codes in order to stress the necessity of existing configurations. New or alternative values and demands appear as the introduction of ideology and inefficiency to those whose requirements are already represented by existing designs.

As Feenberg points out, proposed new designs are compared to stable technologies in terms of cost and efficiency and a trade-off is often formulated between the existing, efficient and stable technology and the new, costlier and seemingly inefficient design that attempts to satisfy different requirements (p. 94). Although there may be transitional costs involved in moving from one design to another, Feenberg argues that these costs are just that. Just like their predecessors, new designs can be optimised over time and successively improved. Once the very definition of the technology has been redefined and internalised by the technological regime, external values and demands that were once deemed as unnecessary and ideological become unquestioned and define the technology.

Agency, Locality, and Representation

Feenberg suggests that agency, representation, and locality need to be reconceptualized for the technical sphere, since traditional political formulations have difficulty accounting for democratic interventions into technical design (pp. 105-106). This is not to say that traditional forms of voting and regulation are not important or that they do not play a valuable role, but merely to point out that they cannot be solely relied upon to guarantee that harmful designs will be avoided or that traditional forms of representation and locality based on geopolitical boundaries can be mobilized around technical issues. Feenberg's alternative attempts to account for democratic interventions into the technical sphere and is formulated in terms of technical codes and a third symmetry intended to supplement the symmetry of successful and unsuccessful technologies and the generalized principle of symmetry introduced by actor network theory.

This third symmetry, which Feenberg calls the symmetry of programs and anti-programs, extends Latour's (1992) "program of action," which may be understood as an attempt to achieve particular outcomes or effects through a distribution of action among human and non-human actors. In order to be successful at implementing programs, the network builder needs to associate and fix actors in such a way that they resist disassociation and contribute to the achievement of the program. Understood in these terms, the "anti-program" is conceived as the disaggregating forces and tendencies of indifferent or unhelpful actors that the network builder must overcome in order to be successful.

Feenberg suggests that the anti-program should not be understood simply as a negation, but instead should be considered as an alternative program (Feenberg, 1999, p. 119). That is, if the network can be taken up and reconfigured by actors incompletely enrolled or fixed by the original program then this effort to reconfigure the network appears as a disaggregating tendency only from the viewpoint of the original network builder. What is needed is an account that symmetrically considers the programs and activities of all actors. This third symmetry, together with an understanding of technical codes, provides the basis for a reformulation of representation, locality, and agency in the technical sphere.

In contrast to traditional forms of geopolitical representation, technical representation is bound temporally by means of technical designs and disciplines that embody a heritage of past interests and choices. As Feenberg states, "technology is the bearer of a tradition that favors specific interests and specific ideas about the good life" (p. 139). This "technical historicity" is materialized in the technical codes of specialized disciplines and is typically invisible to both users and specialists who operate with a seemingly pure and rational autonomy.

In terms of locality, Feenberg suggests that we consider the technical "global" to refer to the larger networks in which the "local" corresponds to the institutional settings in which resistances emerge (p. 139). Conceptualized as such, locality may include both geographically proximate as well as geographically dispersed political interventions into technology, such as may occur around a local site of pollution or around the side effects of a medication experienced by dispersed patients from around the globe. The technical network itself serves as the site of organization and protest. For Feenberg, where individuals act in these "local" settings, they re-enact a form of populist participation traditionally associated with local geographic settings.

Lastly, Feenberg introduces the notion of "participant interests" to refer to the interests that tie individuals to these "locales" (p. 140). That is, individuals enrolled in networks, whether intentionally or unintentionally, have a stake in the impacts of technical activity and, hence, its design and configuration. For example, workers have participant interests in such things as skill levels and safety whenever new technologies are introduced into the workplace. These impacts vary from technology to technology and include the requirements imposed on people and the environment, unintended side effects, as well as other benefits and costs associated with technical activity.

Research Questions

Given the key role played by the Internet Protocol in defining the limits and boundaries of what is and is not possible on the Internet, this thesis seeks to answer the following research questions. First, who was involved in the design of the next generation Internet Protocol (IPng) and what did these participants understand as problems? That is,

- Which relevant social groups were involved in the IPng design?
- Why did these groups want to construct this technology?

Second, and based in part on the answers to the above questions, who was notably absent from the IPng design process and which issues were never articulated? More specifically,

- Who were the irrelevant social groups?
- What participant interests did not inform the design of IPng?

Third, what were the strategies of the involved groups and how successful were these groups in ensuring that their participant interests were accommodated by the eventual IPng design? That is,

- Who/what did the actors interest and enrol?
- What were the various programs and anti-programs?
- What technical codes were produced as a result of the IPng design process?

Last but not least, how did the technological regime of the Internet contribute to the definition of IPng? Specifically,

- What technical codes influenced the IPng design process?

Taken together, these research questions enable us to analyze the actors involved and the strategies employed during the IPng design process. An investigation of the above questions necessarily results in a formulation that outlines the actors involved, their characteristics, interests, and power relations as well as their relation to the Internet Protocol. Consequently, this formulation enables us to deduce whom the design process excluded, the problems that were not articulated, and the values and interests that were not represented in the final design. Lastly, answers to these questions enable us to evaluate the affordances and limits of the next generation Internet architecture and, consequently, the boundaries of new media that rely on this architecture.

CHAPTER 3: METHODOLOGY

Introduction

In this chapter I describe the methodological strategy adopted for this thesis. This strategy is informed by the research questions and the ontological and epistemological position outlined in the previous chapter. Specifically, I summarize the requirements and features of a methodology that is capable of registering the complexity, richness, and heterogeneity of actor-networks and, consequently, that is well suited to an examination of the research questions that guide this thesis. To this end, I rely mainly on Latour (1999) and describe in further detail the concepts of actor, network, and actor-network. I also outline the status and role of the text in terms of its relation to the actor-network as well as with respect to description and explanation. Given the shortcomings of constructivist approaches outlined in the previous chapter, I suggest that a methodology that combines insights from actor network theory and the critical theory of technology is well suited to the research questions that guide this thesis.

I devote the remainder of the chapter to describe in detail and to justify the textual analysis conducted for this thesis. I begin with an overview of the Internet standardization process and locate the Request for Comments (RFC) documents within this process. I argue that RFCs published with respect to IPng, understood in terms of trials of strength and weakness, enable us to follow the actor-network of IPng and the programs and anti-programs of the involved participants. I outline the criteria and procedures used to select the RFCs for analysis and describe the interpretative framework used to examine these texts. I conclude the chapter with a discussion of the validity, reliability, and generalizability of the conducted research.

Actors, Networks, and Actor-Networks

Actor network theory places the responsibility of "theory" on a methodology that is capable of accounting for the detail, complexity, and heterogeneity of actor-networks (Latour, 1999). That is, by avoiding assumptions about how actors could or should behave and how they can and cannot be related, actor network theory opens up an empty methodological frame against all a priori assumptions in order to make possible the recording of actors and their associations. As Latour argues, this is not to say that anything goes or that anything is possible with actual actor-networks, but simply to stop the researcher from prescribing the limits, possibilities, and shapes beforehand.

The notion of network is used to indicate a particular topology that features interconnected nodes (Latour, 1999). This topology is clearly different from a three-dimensional Euclidean topology, where objects occupy a volume within a larger space. The topology of the network allows objects to have as many dimensions as they have connections. There are several important features and implications that are made possible by and that result from this choice of topology.

First, the approach avoids a priori commitments to proximity, scale, order and all other specific types of relations (Latour, 1999). That is, distinctions between far/close, micro/macro, top/bottom, inside/outside are understood as effects or outcomes that must be produced and maintained. The notion of network dissolves all of these a priori distinctions and shifts focus to the number, type, and topography of connections.

Second, strength and durability are understood as features that need to be achieved and secured by connecting, relating, and tying elements together (Latour, 1999). Strength is achieved from the weaving of ties that are themselves woven from weaker ties. That is, strength is achieved through heterogeneity, not homogeneity.

The notion of actor, like the notion of network, has a particular definition in actor network approaches. Put simply, an actor, or "actant," is anything that acts or that is granted activity by others (Latour, 1999). This definition does not limit actors to individual human actors and is purposefully broad enough to encompass a full range of humans and non-humans. The only requirement that every actor needs to satisfy is one of action.

Text, Description, and Explanation

By moving from a semiotics of meaning to a semiotics of things, actor network theory dissolves the distinction between meaning production and things "out there" – a distinction typically assumed by semioticians accustomed to studying, for example, fiction or popular culture (Latour, 1999). The key insight gained by applying semiotics to a study of science and technology is that discourses and texts also define the things assumed to be "out there." Texts, discourses and narratives are granted the power to define their social context, the nature assumed to be "out there," their authors, their readers, as well as themselves.

Understood in this way, texts are a means of attributing characteristics and competencies to actors, distributing actions across actors, and establishing connections and relations between actors (Latour, 1999). Like any narrative, attributions, distributions, and relations are understood as produced and generated. Similarly, like textual characters, actants are not conceived as fixed entities, but instead are understood as fluid, circulating objects that are undergoing trials of strength and actions that support or undermine their stability, durability and continuity. Again, this is not to say that actual actors are infinitely pliable and flexible. Certain actors are in fact quite stable and durable. However, this durability and stability is always understood as something that is maintained and achieved. Durable actors are those actors that are able to withstand trials of strength that attempt to disassociate and destabilize them.

In terms of description and explanation, the actor-network traces itself as it changes, grows, and expands (Latour, 1999). The process of binding unrelated elements, attributing characteristics and undergoing trials produces explanatory resources that cannot be separated from the actor-network. Every actor-network is inseparable from its own formation and change and from its own frame of reference. In this sense, description and explanation are one and the same. As Latour (1999) points out, the actor-network is not some thing, but the recorded movement of a thing. The explanation lies in the description of how unrelated elements were made to be related, how competences were assigned, and how actions were distributed.

Strictly speaking, the methodology requires that the researcher not add causes from outside of the actor-network, since this would imply that they are somehow connected to the network. However, as discussed in the previous chapter, this requirement is problematic, since explanations that simply describe traced actor-networks risk normative affirmations of the existing state of affairs (Feenberg, 1999, p. 31). This criticism implies that a critical approach must include an evaluative component that explicitly takes into account technology's ability to affirm and undermine political, economic, and social relations. Consequently, the methodology undertaken for this thesis may be understood as composed of two parts. First, a strict methodological frame is maintained to record the various actor-networks that constitute the IPng design process. Second, the critical theory of technology is used as a basis to evaluate the traced actor-networks and to reflect on their relation to wider political processes.

Internet Standardization

In order to be able to follow the attribution of characteristics and the distribution of action with respect to IPng, we need to now turn to the Internet standardization process, which each technology specification must successfully traverse in order to be considered an Internet standard. Although the Internet standardization process has undergone many changes over the years, it continues to provide the participants involved with a detailed structure for subjecting new specifications to trials of strength and weakness. In fact, the Internet standardization process as a whole may be understood as a single, long-term trial of strength that tests and assesses a proposed specification. Specifications that withstand this trial of strength are eventually published as Internet standards. In contrast, specifications that fail this trial of strength are never converted into standards. Their existence is simply marked by publications, documented discussions, and decisions up until the point in the process where the specification was deemed a failure.

When we look at the Internet standardization process in terms of the parts that make it up, we can see that a specification and its proponents are subjected to increasingly difficult trials of strength as the specification proceeds along in the process and overcomes previous trials of strength. More specifically, when the first public discussions began to take place regarding IPng the Internet standardization process involved the following procedures:

1. A specification is published by an individual, organization, or Internet Engineering Task Force Working Group (IETF WG) (Chapin, 1992, p. 8).
    - In most cases, a specification that originates outside of an IETF WG is submitted to an appropriate IETF WG for review and revision. If an appropriate IETF WG does not exist a new group is created (p. 8).
2. An initial draft version of the specification, known as an Internet Draft, is made available to the public for review, revision, and comment (p. 5).
    - An Internet Draft must remain available to the public for a minimum of 2 weeks before it is submitted to the Internet Engineering Steering Group (IESG) (p. 5).
3. Within the IESG, the appropriate IETF Area Director, or the Chairman of the IETF, reviews the Internet Draft (p. 7).
    - In consultation with the IESG, the IETF Area Director or IETF Chairman may commission an independent technical review of the specification (p. 7).
    - If the specification has broad impact for the Internet, the IESG may form and commission a special review and analysis committee to evaluate the specification (p. 7).
    - If the specification has broad impact for the Internet and the criteria for advancement along the standards track are not universally recognized, the IESG may commission the development and publication of criteria that need to be met by the specification (p. 7).
    - Upon completion of its evaluation, the IESG communicates its findings to the IETF community (p. 7).
4. The IETF community reviews the findings of the IESG and has an opportunity to raise any issues that it feels are significant and need to be resolved by the IESG (p. 7).

5. Based on IETF community feedback, the IESG makes any necessary revisions and decides whether the specification should enter the standards track. It communicates its decision to the Internet Architecture Board (IAB) as a recommendation for action (p. 7).
6. The IAB reviews the IESG's recommendation for action. If it finds any significant problems with either the decision or the specification, it tries to resolve them with the Working Group and its chair and/or the document author(s), with the assistance of the IESG and the relevant IETF Area Director (p. 7).
7. If the IAB accepts the IESG's recommendation, the specification becomes a "Proposed Standard" and is published as a "Standards Track" RFC (p. 8).
8. In order for a specification to move from "Proposed Standard" to "Draft Standard," and then from "Draft Standard" to "Standard," steps 3 – 7 need to be repeated for each attempted movement toward "Standard" designation (pp. 7, 11-12).
    - In step 3, the specification is not an Internet Draft but a published RFC.
    - In step 5, the IESG does not decide if the standard should enter the standards track, but whether it should advance along it.
    - In step 7, the specification is designated either as "Draft Standard" or "Standard."
9. If a specification has not reached "Standard" designation within 24 months of entering the standards track, the IESG reviews the specification and makes a recommendation as to whether the specification should be terminated or whether it should continue along the standards track. The IESG communicates its recommendation to the IETF community (p. 9).
10. Based on IETF community feedback, the IESG makes any necessary revisions to its recommendation, which it communicates to the IAB (p. 9).
    - Based on the IESG recommendation, the IAB makes the final decision on whether to terminate the specification or whether to let it continue in the standards track (p. 9).

As the above description indicates, what initially enters the standardization process as an Internet Draft is continually and repeatedly subjected to increasingly difficult trials of strength. IETF Working Groups, organizations, review and analysis committees, IETF Area Directors, and IAB members place the specification under strain in order to determine whether it can perform as an Internet standard. A specification that receives "Proposed Standard" designation is "generally stable, has resolved known design choices, is believed to be well-understood, has received significant community review, and appears to enjoy enough community interest to be considered valuable" (p. 11).

If a specification is able to "survive," receive "Proposed Standard" designation, and reach the latter stages of the Internet standardization process, it still has to face its most difficult trials.

Consider the criteria and requirements that must be met by specifications at this stage in the process.

- "Proposed Standard":
    - A specification must remain at the "Proposed Standard" designation for at least 6 months (p. 8).
    - Although implementation and operation experience is not required for "Proposed Standard" designation, such experience strengthens a specification proposal (p. 11).
    - A specification that has broad impact on much of the Internet requires implementation and operation experience before it will be considered for "Proposed Standard" designation (p. 11).
    - If major changes are needed in order to advance a specification from "Proposed Standard" to "Draft Standard" the IESG may recommend that the specification remain at its current designation in order to mature and develop (pp. 8-9).
- "Draft Standard":
    - A specification must remain at the "Draft Standard" designation for at least 4 months (p. 11).
    - If major changes are needed in order to advance a specification from "Draft Standard" to "Standard" the IESG may recommend that the specification remain at its current designation in order to mature and develop (pp. 8-9).
    - In order to be considered for the "Draft Standard" designation, a specification must have "adequate operational experience" and at least two independent implementations that are able to interoperate with each other (p. 12).
    - In order to be a "Draft Standard," a specification "must be well-understood and known to be quite stable, both in its semantics and as a basis for developing an implementation" (p. 12).
- "Standard":
    - If major changes are needed in order to advance a specification from "Draft Standard" to "Standard" the IESG may recommend that the specification re-enter the standard track as a "Proposed Standard" (pp. 8-9).
    - A specification is elevated to "Standard" designation when "significant implementation and operational experience has been obtained" (p. 12).
    - Unlike "Draft Standard" specifications, "Standard" specifications have been implemented and tested in large-scale production environments (p. 12).

At a minimum, a specification that is designated "Proposed Standard" must undergo testing, review, and revision for an additional 10 months before it can be considered for "Standard" status. During this time, at least two independent implementations will need to be made to interoperate with each other. Moreover, implementations will need to be tested on a large-scale, outside of laboratories, such that they can be made to demonstrate that they can perform under stress in production environments. In this way, the standardization trials of strength harden specifications into standards and initial blueprints into working and performing technologies.

This movement through trials of strength from initial Internet Drafts to proven implementations may be understood in terms of Latour's (1987, p. 138) sociogram and technogram and Feenberg's (1999, p. 88) technical code. That is, the trials of strength that specifications are made to undergo at the hands of relevant social groups mould and shape the specification in such a way that relevant social groups need to be satisfied if the specification is to move toward closure. Unsuccessful specifications die at the hands of an unsatisfied social group and a trial of strength that the specification is unable to overcome. Specification reviews, revisions, changes, and translations reflect the shaping that specifications undergo in order to keep the relevant social groups interested. In this way, the needs, requirements, and interests of relevant social groups become embedded in the specification as technical codes.

In addition, non-human elements such as implementations, network test-beds, and experiment results are associated with specifications to enrol and control the actors necessary to turn a specification into a standard. All specification authors, if they are to succeed in developing a standard, need to enrol the IESG and the IAB by means of these non-human elements. Specification authors unable to produce, control, and associate these non-human elements with their specifications simply cannot control the IESG and the IAB, who eventually lose interest and terminate the specification.

Internet standardization is also inseparable from the various texts that are required during different stages of the process. It is here, via these various texts, that we can observe the attribution of characteristics, the distribution of action, and the association of a specification's constituent elements. More specifically, the Internet standardization process includes the following types of text, understood in the broadest sense of the term to include both the written and spoken word as well as all forms of visual depictions:

- Internet Drafts
- IETF meeting presentations, discussions, and minutes
- IETF mailing lists
- Request for Comments publications

For the purposes of this thesis, the Request for Comments publications were selected for analysis for several reasons. First, and most importantly, the RFCs published by IETF Working Groups may be understood as presenting a concise summary and position of the Working Group. That is, members of IETF WGs discuss, debate, and review their own work through meetings, presentations, and mailing list discussions. In other words, IETF WGs subject their own work to trials of strength and weakness and produce various drafts and working copy documents. Upon reaching some form of consensus, certain drafts and documents are submitted for publication as RFCs to report on the Working Group's activities and findings. As a result, RFCs published by IETF WGs trace the actors and the associations that have withstood the many trials of strength and weakness from the Groups' meetings, mailing list discussions, and debates.

In a similar way, organizations and groups other than IETF WGs submit RFCs for publication in order to present and outline a position, preference, or concern with respect to proposed specifications and standards. These RFCs, like those published by IETF WGs, explicitly purport to summarize and present the position of the organization or group in question. Consequently, they can also be understood in terms of the trials of strength, internal to the group or organization, that have given shape to the actor-network traced by the RFC.

This is not to say that there are no trials of strength remaining following the publication of an RFC. On the contrary, additional trials of strength remain. However, these trials of strength must now stress and strain actors and associations that have been made somewhat durable by previous trials of strength resulting from discussions, revisions, and meetings. As evidenced by the detailed procedures of the standardization process, RFCs published with respect to a proposed specification include RFCs published by IETF WGs, interested groups and organizations, as well as the IESG, which at times publishes criteria requirements as well as its own evaluation of a particular specification. RFCs, like journal articles, reference each other and in the process support, weaken, extend, and modify themselves. An analysis of RFCs enables us to follow the relevant social groups, actors, translations, and trials of strength, but at a different stage than the meetings, mailing list discussions, and working drafts that preceded their publication.

An analysis of RFCs also requires less time and resources than a comparable analysis that also examines mailing list discussions, meeting minutes, and document drafts. The price of not analysing these additional texts involves a loss of detail and granularity with respect to the involved participants, their interests, and the various strategies for strengthening their own positions, translating the interests of others, and weakening the positions of adversaries. However, I believe that the distinction outlined above warrants the exclusion of these additional data sources and makes available additional time and resources for the task at hand.

Selection Criteria and Procedure

A total of 42 RFCs were selected for analysis. These RFCs were selected using two separate methods. First, the official online RFC repository (RFC Index Search Engine, n.d.) was searched using the following search criteria:

- Search for the term "ipng"
- Search all fields (number, title, author, keywords)
- Search all RFC documents
- Search by means of an entire word match

The search produced 41 RFCs. Since this search did not include RFC abstracts in its search criteria, an additional, manual search was conducted. The same text, "ipng," case insensitive, entire word match, was used. All of the "Summary" RFCs, which include the titles and abstracts of other RFCs, were searched. With the exception of RFC 800, the Summary RFCs are numbered using "99" for the last 2 digits and include all of the RFCs from RFC 699 to RFC

3599, with the exclusion of RFC 799 and RFC 3399, which are not "Summary" RFCs. This manual search produced one additional RFC: *RFC 1681: On Many Addresses per Host*. This RFC, in addition to the 41 RFCs found using the online database search yielded the following set of documents:

- *RFC 1454: Comparison of Proposals for Next Version of IP*
- *RFC 1475: TP/IX: The Next Internet*
- *RFC 1550: IP: Next Generation (IPng) White Paper Solicitation*
- *RFC 1621: Pip Near-term Architecture*
- *RFC 1622: Pip Header Processing*
- *RFC 1667: Modeling and Simulation Requirements for IPng*
- *RFC 1668: Unified Routing Requirements for IPng*
- *RFC 1669: Market Viability as a IPng Criteria*
- *RFC 1670: Input to IPng Engineering Considerations*
- *RFC 1671: IPng White Paper on Transition and Other Considerations*
- *RFC 1672: Accounting Requirements for IPng*
- *RFC 1673: Electric Power Research Institute Comments on IPng*
- *RFC 1674: A Cellular Industry View of IPng*
- *RFC 1675: Security Concerns for IPng*
- *RFC 1676: INFN Requirements for an IPng*
- *RFC 1677: Tactical Radio Frequency Communication Requirements for IPng*
- *RFC 1678: IPng Requirements of Large Corporate Networks*
- *RFC 1679: HPN Working Group Input to the IPng Requirements Solicitation*
- *RFC 1680: IPng Support for ATM Services*
- *RFC 1681: On Many Addresses per Host*
- *RFC 1682: IPng BSD Host Implementation Analysis*
- *RFC 1683: Multiprotocol Interoperability In IPng*
- *RFC 1686: IPng Requirements: A Cable Television Industry Viewpoint*
- *RFC 1687: A Large Corporate User's View of IPng*
- *RFC 1688: IPng Mobility Considerations*
- *RFC 1705: Six Virtual Inches to the Left: The Problem with IPng*
- *RFC 1707: CATNIP: Common Architecture for the Internet*
- *RFC 1710: Simple Internet Protocol Plus White Paper*
- *RFC 1715: The H Ratio for Address Assignment Efficiency*
- *RFC 1719: A Direction for IPng*
- *RFC 1726: Technical Criteria for Choosing IP The Next Generation (IPng)*
- *RFC 1752: The Recommendation for the IP Next Generation Protocol*
- *RFC 1753: IPng Technical Requirements Of the Nimrod Routing and Addressing Architecture*
- *RFC 1776: The Address is the Message*
- *RFC 1883: Internet Protocol, Version 6 (IPv6) Specification*
- *RFC 1884: IP Version 6 Addressing Architecture*
- *RFC 1885: Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6)*
- *RFC 1886: DNS Extensions to support IP version 6*

- *RFC 1887: An Architecture for IPv6 Unicast Address Allocation*
- *RFC 1955: New Scheme for Internet Routing and Addressing (ENCAPS) for IPNG*
- *RFC 2460: Internet Protocol, Version 6 (IPv6) Specification*
- *RFC 3194: The H-Density Ratio for Address Assignment Efficiency An Update on the H ratio*

As can be seen from the above titles, the selected sample includes a wide variety of RFCs published by various individuals, groups, and organizations. Included are RFCs published by organizations concerned with the requirements of the new protocol, such as *RFC 1669: Market Viability as a IPng Criteria*, and *RFC 1674: A Cellular Industry View of IPng*. Also are included RFCs published by IPng Area Directors, such as *RFC 1550: IP: Next Generation (IPng) White Paper Solicitation*, and *RFC 1752: The Recommendation for the IP Next Generation Protocol.* Based on this initial overview and the previously outlined and exhaustive, two-step method to locate RFCs related to IPng, the above sample was deemed adequate and well suited to the research questions that guide this thesis.

Interpretative Framework

Based on the research questions outlined in the previous chapter, each RFC was analysed in terms of the following concepts: relevant social group, participant interest, technical code, program/anti-program, actant, translation, and trial of strength. Integrated together, and applied across all of the selected RFCs, this interpretative framework forms the basis from which the research questions will be addressed. The following sections define more precisely each of the above terms.

Relevant Social Group

A relevant social group is a collection of individuals who share the same set of meanings with respect to a particular technology (Pinch & Bijker, 1987). Relevant social groups may include unorganized groups of individuals as well as organized groups such as institutions and organizations.

Problems, Solutions, and Participant Interests

Relevant social groups constitute problems for which the technology may provide potential solutions (Pinch & Bijker, 1987). These explicit problems may also be understood as the explicit interests, or the participant interests (Feenberg, 1999, p. 140) of relevant social groups involved in sociotechnical networks. Different groups may propose different technical solutions based on their understanding of the technology and the potential problems that it may address. Different technical solutions may also be proposed by different members of a particular relevant social group.

Technical Code

Technological regimes, which encompass scientific knowledge and institutions, engineering practices and procedures, and standardized definitions of technologies, conceal the social and cultural contingencies of technologies in a seemingly neutral and highly technical language. The notion of technical code (Feenberg, 1999, p. 88) foregrounds the relation between participant interests, sociotechnical conflict and what over time appear as self-evident and seemingly necessary aspects of technologies and technological regimes.

Programs and Anti-programs

The problem and solution claims of relevant social groups may lead to a variety of conflicts. These conflicts may include conflicts over technical requirements that address certain concerns at the detriment of others as well as conflicting solutions to the same problem (Pinch & Bijker, 1987, p. 35). Understood in terms of Feenberg's (1999, p. 119) symmetry between programs and anti-programs and technical locality (1999, p. 139), these conflicts are manifested in the attempts of various actors to associate human and non-human elements in such a way that the resulting association resists the disaggregating forces of the enrolled actors. To the extent that this enrollment may include other relevant social groups, conflicts may arise over the configuration of the actor-network such that these social groups may introduce new elements and associations or modify existing elements and associations in an attempt to stabilize new or alternative sociotechnical configurations.

Actant

An actant, or actor, is anything that acts or that is granted activity by others (Latour, 1999). The notion of actant is not limited to individual human actors, but instead encompasses a full range of human and non-human actors.

Translation

Translation is the interpretation given by relevant social groups of their own explicit interests and that of the actors that they try to enrol to further their efforts (Latour, 1987, p. 108). Translation is a method for associating otherwise unrelated actors. It is also a means of modifying or doing away with the explicit interests of relevant social groups involved in the construction of a technology.

Trials of Strength

Trials of strength and weakness determine what is "real" and what is not "real" (Latour, 1988, pp. 158-159). Understood in this way, the real is not one thing among others that are unreal, but instead is understood in terms of a spectrum of resistance. A shape is the outline that results from trials of strength and weakness, which deform, transform, and perform it.

Validity, Reliability, and Generalizability

Questions regarding the quality of "qualitative" research are typically problematic when the terms and concepts used to judge research designs are imported directly from "quantitative" disciplines. Although a detailed overview of the debates and problems in this area is beyond the scope of this thesis, I agree with Mason (2002, p. 38) that issues of quality need to be addressed in all forms of research, regardless of their "qualitative" or "quantitative" designation. To this end, I followed Mason's orientation with respect to validity, reliability, and generalizability and took into account the practical strategies entailed by this orientation with the explicit goal of producing quality research.

According to Mason (2002, p. 38), research is valid if the researcher observes, identifies, or "measures" what he or she claims to be observing, identifying, or "measuring." The key activity is to demonstrate that concepts stemming from the ontological and epistemological perspective of the research may be identified via the selected method(s) and data sample(s). This requirement is satisfied by the research conducted for this thesis by means of an explicit relation between the ontology, epistemology, and methodology underlying the research activity. That is, as I have outlined in this and the previous chapter, texts are not understood as hovering above and beyond some other reality, which they may describe more or less accurately. Instead, reality is understood to be simultaneously real, discursive, and social. Texts are understood to be part of and constituent of this reality, which they trace for the researcher. The concepts introduced in this chapter provide a minimalist interpretative framework for the reading of the selected texts and are intended to foreground particular elements of traced actor-networks without prescribing any unnecessary limits on how exactly these elements may be traced.

As Mason (2002, p. 39) points out, research reliability in "quantitative" disciplines is typically achieved through double-checking of results yielded by particular methods and/or by means of cross-checking results using different methods. In these disciplines, the key question is whether the instruments, tools, and methods reliably and accurately assess the data and lead to repeatable results. For the purposes of this study, the notion of reliability needs to be conceptualised in such a way that it is meaningful to talk about "accuracy" and "repeatability."

Unlike other discourse analyses, which often make use of a complex metalanguage to classify and organize the reading and interpretation of texts, the analysis conducted for this thesis uses a simple and minimalist language for interpretation. This decision is motivated by and follows from a methodological commitment to reduce as much as possible a priori classification and prescription of what may and may not be observed. A classificatory or interpretative framework based on language structure or function is incompatible with these goals, since to some extent it prescribes in advance the sort of language patterns and uses that are possible given particular contexts.

Giving up on a structured framework does come at a price. The minimalist framework adopted for this thesis does not specify how the various concepts can be expected to appear in the text. Consequently, there is much room for interpretation in terms of what exactly counts as a trial of strength, a translation, or a technical code, for example. To minimize this variability and increase as much as possible the accuracy and reliability of the analysis, the following practices were followed:

- The criteria and justification for text selection was described in detail.
- The selected texts were enumerated and are widely available for review.
- Findings are supported by means of direct quotations and text extracts.

With respect to the generalizability of the findings, or the extent to which wider claims may be made based on the research conducted for this thesis (Mason, 2002, p. 39), I make only three small claims. First, that the findings of this research may resemble in some respect the circumstances of another technology in the making. In this case, the findings may be used for comparison and further evaluation. Second, the findings of this research may be instructive for future technology in the making, in particular future protocols and standards. In this case, the findings may be instructive in terms of future strategies and approaches to standardization or technology construction. Lastly, and somewhat related to the first point, the findings of this research may inform existing theories of technology by foregrounding elements that were previously overlooked or incompletely theorized.

CHAPTER 4: FINDINGS

Introduction

In this chapter I answer the following research questions: Which relevant social groups were involved in the design of IPng and what did these groups understand as problems? What programs and anti-programs were involved? And lastly, what technical codes were produced as a result of the IPng design process? To answer these questions, I begin with an introduction of the IETF IPng Area, the relevant social group responsible for the design of IPng. I suggest that by means of a solicitation strategy the IETF IPng Area interested other social groups necessary to the eventual adoption of IPng. By allowing these groups to contribute to the definition of IPng, I contend that the IPng Area was able to design an IPng that is well designed to enrol these groups in the future.

For each relevant social group that responded to the IPng Area's solicitation, I outline in detail the actor-network traced by the group, foreground the group's participant interests, and show how these participant interests were translated into particular technical codes. I conclude the section with a discussion of the IPng Area's inclusion of these technical codes in a technical criteria document used to evaluate actual IPng specification proposals. I summarize in detail the actor-network traced by the authors of this criteria document and suggest that it acknowledges the groups' participant interests through the inclusion of technical codes that are deemed necessary for all IPng proposals.

I end the chapter with an analysis of the IPng specification that was eventually selected and adopted by the IETF IPng Area. I suggest that it is largely based on the technical codes of the relevant social groups and, as a result, is well suited to their enrollment. I also conclude that IPng was explicitly designed to accommodate the widest possible variety of groups and interests. In particular, through the use of extensible header options, IPng is designed to avoid many potential design conflicts that may arise from competing programs.

The IPng Area and RFC 1550: Interesting and Enrolling Others

The directorate of the IETF IPng Area includes directors who "are experts in security, routing, the needs of large users, end system manufacturers, Unix and non-Unix platforms, router manufacturers, theoretical researchers, protocol architecture, and the operating regional, national, and international networks" (Bradner & Mankin, 1995, p. 6). The IPng Area cites a dwindling IPv4 address space and growing routing tables as the primary motivators for improving IPv4 (pp. 6-7). Although at first the IPng Area is unsure whether there is sufficient time to add new functionality to IPv4 in addition to including support for more addresses, based on projections of expected address depletion, a decision is made to consider additional functionality as part of IPng proposals (p. 7).

In order to determine the scope and direction for this additional functionality, the IPng Area published RFC 1550, *IP: Next Generation (IPng) White Paper Solicitation* (Bradner & Mankin, 1993). RFC 1550 invites all "interested parties" to submit papers "detailing any specific requirements that they feel an IPng must fulfill or any factors that they feel might sway the IPng selection" (p. 1). That is, according to RFC 1550, submitted documents will inform the IPng design and decision making process and "will be used as resource material by the various IPng working groups, the directorate, the external review board and the area directors" (p. 3).

Understood in terms of Latour's (1987) sociogram and technogram, this solicitation may be conceived as a strategy of interest and enrollment. That is, in order for IPng to be successful as a technology, it must eventually be taken up by various social groups and used as an unproblematic black-box. The IPng Area, with its working groups, directors, and review boards, can improve the likelihood of this adoption by letting the various social groups translate their interests into IPng design features such that when these features are implemented the IPng Area can use them to enrol the groups. A deficiency in the sociogram, that is, uninterested or indifferent social groups, is overcome through the association of non-human elements, IPng design features in particular, in the technogram. What is unique in this case is that the actors that the IPng Area wishes to enrol are given the opportunity to contribute to the definition of the technogram, and consequently, to their own enrollment. Generally speaking, this strategy is similar to participatory and user-centric approaches to technology design and construction.

Documents submitted in response to the IPng Area's solicitation were reviewed by members of the IPng Area to clarify any ambiguities and to assess technical feasibility (Bradner & Mankin, 1993, p. 2). Following this review and any revisions by authors, submitted documents were submitted as Internet Drafts to the IETF community for comment. Unless withdrawn by the author(s), the Internet Drafts were then submitted for publication as informational RFCs (p. 2). Based on an examination of these RFCs, 21 relevant social groups were identified in addition to the IETF IPng Area.

The Interested Parties

In this section I describe in detail the relevant social groups that responded to RFC 1550. The description of each relevant social group includes the following details. First, I briefly introduce the relevant social group on its own terms using the group's self-description. Second, I describe the actor-network traced by the author(s) that claim to represent each relevant social group. This description outlines the various translations that associate otherwise separate elements and includes an identification of the group's participant interests and the proposed solutions that correspond to these interests. Taken as a whole, the description may be understood as the group's program. Each group translates IPng in terms of this program such that IPng is attributed certain characteristics and features that enable its particular association with the actor-network. As the descriptions illustrate, the general strategy of each group is to trace an actor-network that includes various durable and stable elements which IPng must be capable of accommodating.

I conclude each description with a summary of the technical codes that are implied by the group's program. For each summary, I use simple descriptions that sum up the technical features required of IPng by each relevant social group. Although I use a consistent and descriptive naming scheme across relevant social groups, this should not be interpreted to imply that technical code descriptions map to identical IPng characteristics for each relevant social group. Although there are overlaps across relevant social groups in terms of IPng requirements, the exact definition of the technical code description is always obtained through reference to the detailed description of the group's program. The technical code summary is an attempt to account for general similarities and differences across relevant social groups and is intended as a guide to the reader.

With respect to the assignment of technical code descriptions, the following conventions were followed. First, I distinguished between direct and indirect support for a feature. For example, in the case of security, direct support by IPng could involve the encryption of packet data at the network layer. In this case, the technical code description would be, "Security support, direct." In contrast, if a relevant social group simply required the ability to secure data and was content on using upper layer protocols to achieve this, the technical code description would be, "Security support." That is, IPng should not prohibit the implementation of security mechanisms at other layers, but does not itself require direct support for security mechanisms.

In addition, technical code summary descriptions were assigned in accordance with the level of detail provided by each group. For example, one group may require that IPng support high performance subnetwork technology but may not specify in detail which technologies in particular this may include, whereas another group may also require this support but may also specify the exact technology or technologies in question. In the case of the former, the technical code summary description would be, "High performance subnetwork support," whereas in the latter it could be something such as, "ATM support." Although at first glance it may seem as though the two descriptions imply each other, it may be the case that a relevant social group only requires support for ATM and is completely indifferent to other technologies. Likewise, a group may require support for high performance subnetwork technologies, but is indifferent or undecided with respect to ATM specifically. Rather than decide and arbitrate such cases myself, I left the ambiguity intact and distinguished these cases according to the detail provided or withheld by each relevant social group.

To guide the reader, the relevant social groups are described in the order below. The names of the groups follow as closely as possible either the title of the RFC that was published by each respective group in response to RFC 1550 or some distinguishing feature of the group's self-description.

- A. The Defense Modeling and Simulation Community
- B. Unified Routing Architecture
- C. Market Viability
- D. IPng Engineering
- E. IPng Transition
- F. IPng Accounting
- G. Electric Power Research Institute
- H. Cellular Digital Packet Data Consortium of Service Providers
- I. IPng Security
- J. Italian National Institute for Nuclear Physics Network Team
- K. U.S. Navy Radio Frequency Tactical Systems Group
- L. Corporations #1
- M. U.S. Navy High Performance Network Group
- N. Asynchronous Transfer Mode
- O. Multiprotocol Interoperability In IPng
- P. The Cable Television Industry
- Q. Corporations #2
- R. IETF Mobile IP Working Group
- S. Multiple Addresses Per Host
- T. Transport and Network Layer Independence
- U. IETF Nimrod Working Group

A. The Defense Modeling and Simulation Community

The Defense Modeling and Simulation community conducts Distributed Interactive Simulation (DIS) battle and warfare exercises using the Internet (Symington, Wood, & Pullen, 1994).

Actor-Network, Participant Interests, and Proposed Solutions

The simulation process relies on multiple and varied networks and hosts communicating with each other via the Internet. These networks and hosts are geographically dispersed around the world (p. 2). Despite these distances, information needs to be transmitted in such a way that the decisions and actions of both human and non-human entities (i.e. enemy commanders, weapon systems, etc.) may be tracked and monitored in real-time, that is, as they could be observed and tracked outside of the simulation (p. 2). The Defense Modeling and Simulation community has already developed and uses Distributed Interactive Simulation protocol and protocol data unit standards, which specify the exact parameters and performance expectations that must be satisfied by simulations (p. 2).

Specifically, the latency between the output of a protocol data unit at the application level of a simulator and input of that protocol data unit at the application level of any other simulator should be "100 milliseconds for exercises containing simulated units whose interactions are tightly coupled" and "300 milliseconds for exercises whose interactions are not tightly coupled" (p. 4). IPng should be able to support these latency requirements. IPng should also provide specific reliability guarantees: "The reliability of the best-effort datagram delivery service supporting DIS should be such that 98% of all datagrams are delivered to all intended destination sites, with missing datagrams randomly distributed" (p. 4).

The data transmitted during a simulation includes what the Modeling and Simulation community considers to be sensitive information, such as the properties of certain weapon systems and warfare strategies, for example (p. 4). The Modeling and Simulation community requires that this information remain confidential. In addition, there is much information that is generated and communicated during the course of a simulation, stressing both the available network bandwidth and the capacity of hosts to handle and process this information (p. 3). Although the Modeling and Simulation community has tried to make use of the existing IPv4 capabilities to deal with these transmission requirements, it has had only limited success given what it considers to be the limitations of IPv4 (p. 4).

IPng should enable multicasting and the management of multicast groups so that hosts can selectively send and receive data. These capabilities should be "[s]calable to hundreds of sites and, potentially, to tens of thousands of simulation platforms" (p. 4). In addition, "group management mechanics must have the characteristics that thousands of multicast groups consisting of tens of thousands of members each can be supported on a given network and that a host should be able to belong to hundreds of multicast groups simultaneously" (p. 5). In terms of management performance requirements, "group members must be able to be added to or removed from groups dynamically, in less than one second, at rates of hundreds of membership changes per second" (p. 6).

IPng should also have resource reservation mechanisms such that certain network traffic may be prioritized over other traffic. For the Defense Modeling and Simulation community "[s]uch a resource reservation capability is essential to optimizing the use of limited network resources, increasing reliability, and decreasing delay and delay variance of priority traffic, especially in cases in which network resources are heavily used" (p. 6).

Technical Code Summary

- Datagram based
- Low latency
- High reliability
- Security support
- Multicasting
- Resource reservation

B. Unified Routing Architecture

Supporters of the Unified Routing Architecture wish to use the Unified Routing Architecture for inter-domain routing on the Internet (Estrin, Li, & Rekhter, 1994).

Actor-Network, Participant Interests, and Proposed Solutions

To provide scalable routing, IPng must provide support for topological address assignment. Since it is difficult to determine in advance how routing information will be aggregated, "the IPng addressing structure should impose as few preconditions as possible on the number of levels in the hierarchy" (p. 1). That is, the number of levels at different parts of the address hierarchy should be allowed to vary. Likewise, the levels should not be statically tied to particular parts of the address. In addition, since the hop-by-hop routing algorithm needs to be able to determine the next hop based on information available in an IPng packet, the "IPng packet format must provide efficient determination of the full hierarchical destination address" (p. 1). However, hierarchical addresses should not imply hierarchical routing, and IPng should "carry enough information to provide forwarding along both hierarchical and non-hierarchical routes" (p. 1).

IPng should also accommodate a routing label that supports at least two types of labels that specify the routing preferences for the packet and define how routers should behave upon encountering each type of label (p. 2). Lastly, IPng should provide support for a source routing mechanism that enables the "[s]pecification of either individual routers or collections of routers as the entities in the source route" (p. 2). As part of this function, IPng should have an "option to indicate that two consecutive entities in a source route must share a common subnet in order for the source route to be valid," as well as the ability to specify "the default behaviour when the route to the next entry in the source route is unavailable" (p. 2). IPng should also support the ability to verify the feasibility of a source route (p. 2).

Technical Code Summary

- Datagram based
- Hierarchical addresses
- Hierarchical routing
- Non-hierarchical routing
- Variable hierarchy levels
- Routing label
- Source routing

C. Market Viability

Based on a review of existing IPng specification proposals, John Curran (1994) argues that existing proposals lack the necessary market viability to be widely adopted and used.

Actor-Network, Participant Interests, and Proposed Solutions

In order for new technologies to succeed in an open marketplace, they must provide consumers with new capabilities and reduced costs (p. 1). Consumer confidence and demand drive the adoption of new technology (p. 1). Internetworking technologies in particular can be difficult to deploy and must provide consumers with a high return on investment if they are to be adopted (p. 1).

Although the adoption of protocols in the past by the computer industry has resulted in general acceptance by the networking industry, there have also been instances where adoption by the computer industry has been insufficient to ensure broader acceptance (p. 1). Even if IPng is adopted by the computer industry, it will need to compete with an established and mature IPv4 (pp. 1-2). Since IPv4 has a large installed base, the computing industry will need to continue to support it (p. 2).

IPv4 consumers will decide to use IPng instead of IPv4 for one of three reasons (pp. 2-3). First, if IPng provides new functionality not found in IPv4. Second, if there are reduced costs from using IPng. And lastly, if consumers need to connect to IPng hosts that they cannot reach via IPv4. With respect to the first requirement, it is not clear whether IPng will provide significant improvements and capabilities over IPv4 (p. 2). With respect to the second requirement, migration to IPng will likely lead to increased costs (p. 2). Finally, in terms of connectivity to

IPng-only hosts, such hosts will only start to appear once all of the IPv4 addresses are depleted. However, customers who are unable to obtain IPv4 network assignment will seek internetworking services and technologies, such as network translation devices, that enable access to the complete Internet instead of settling for services and technologies that can only provide partial access (pp. 2-3).

In sum, IPng must provide new capabilities and significant improvements over IPv4 (pp. 1-3). It also needs mechanisms for transparent access between the IPv4 and IPng communities once the IPv4 address space has been depleted (pp. 1-3).

Technical Code Summary

- New capabilities relative to IPv4
- Reduced cost relative to IPv4
- IPv4/IPng interoperability

D. IPng Engineering

Denise Heagerty (1994) provides an opinion on transition from IPv4 to IPng and other issues that need to be considered if IPng is to succeed.

Actor-Network, Participant Interests, and Proposed Solutions

Personal experience with a DECnet Phase V transition suggests that certain decisions may be made earlier than others when considering protocol transition more generally. Specifically, this experience suggests that administrative changes need the most time whereas routing protocol changes need the least amount of time (pp. 1-2). In order for IPng to succeed, it needs to have a flexible and time sensitive transition plan (p. 1).

IPng decisions and timelines should be broken down into smaller units that are more manageable for service managers (p. 1). In addition, service managers should be provided with a "transition toolbox and scenarios of their uses based on real examples" (p. 2). An effort should also be made to identify essential elements needed for transition so that service managers can transition at their own pace (p. 2). An up-to-date list of necessary software should be maintained and a feedback loop should be used to improve software based on experience (p. 2).

Time is an important factor in the transition process: “A lead time of 10 years (or more) will help to take good long term technical decisions and ease financial and organisational constraints” (p. 2). Service managers need to be able to adapt IPng transition to their unique system environments and their financial, political, and technical conditions (p. 2).

Finally, since certain environments require that systems move between buildings, IPng should not tightly associate an IP address to a physical subnet (p. 2). IPng support for automatic configuration would dissipate this problem (p. 2). In addition, since certain environments need to balance system load across multiple systems, IPng should support hierarchical addressing and routing and should “allow the delegation of network management into subdomains” (p. 3).

Technical Code Summary

- Transition in stages
- Transition autonomy
- Transition management tools
- Automatic configuration
- Hierarchical addressing
- Hierarchical routing
- Decentralized control and policy

E. IPng Transition

Brian Carpenter (1994) provides a personal view with respect to transition between IPv4 and IPng and other general issues that need to be considered in order for IPng to succeed.

Actor-Network, Participant Interests, and Proposed Solutions

The transition from IPv4 to IPng will take many years (p. 1). Network sites will decide when, and whether to make the transition based on their needs and the costs involved (pp. 1-2). Smaller sites may be pressured by Internet service providers and may make the switch in a single step (pp. 1-2). Computer system, router, and application vendors will provide IPng products, but the release of these products will not be coordinated among vendors (p. 2). Users will also continue to use a mixture of newer and older systems and applications (p. 2). Thus, the success of IPng depends on the success of the IPv4 to IPng transition.

As a result, IPng must be designed and implemented in such a way that there is seamless operation between a mixture of IPv4 and IPng hosts and routers. That is, "[a]n IPv4 packet must be able to find its way from any IPv4 host, to any other IPv4 or IPng host, or vice versa, through a mixture of IPv4 and IPng routers, with no (zero, null) modifications to the IPv4 hosts" (p. 2). However, this interoperability should not rely on header translation, since it will create "almost insurmountable practical difficulties" (p. 3). As a result, IPv4 and IPng must be able to coexist on the same hosts and routers (pp. 2-3, 4-5). In order to allow users to manually manage the transition of their systems, IPng should also have support for the manual mapping of IPv4 addresses to IPng addresses (p. 4). Lastly, users are going to need "smart management tools" to manage the transition (p. 5).

There are also a number of groups whose interests need to be satisfied if IPng is to be successful. To begin with, the networking industry invests heavily into new networking technologies and IPng will need to be able to support a variety of underlying technologies (p. 6). Secondly, "[i]t is taken for granted that multicast applications must be supported by IPng" (p. 6). Thirdly, funding agencies want to be able to charge for traffic that flows over their lines such that IPng should include support for policy routing and accounting (p. 6). Lastly, corporate and campus network operators have been the victims of network security violations, which they wish to minimize (p. 7). Consequently, IPng needs to have improved security over IPv4 (p. 7).

Technical Code Summary

- IPv4/IPng interoperability
- IPv4/IPng dual stack per host
- IPv4 to IPng address mapping
- Transition management tools
- Heterogeneous subnetwork support
- Multicasting
- Policy routing
- Accounting support
- Security improvements relative to IPv4

F. IPng Accounting

Nevil Brownlee (1994) suggests that IPng should directly support accounting.

Actor-Network, Participant Interests, and Proposed Solutions

The Internet Accounting Model specifies how accounting information is structured and collected by lower-layer protocols (p. 1). The model can be extended to include IPng and higher-layer protocols (p. 1). Brian Carpenter has suggested one way of tracking policy routing and accounting information via IPng (p. 1). Using Carpenter's method as a starting point, the party responsible, "i.e., willing to pay for, a packet" (p. 2) could be identified. Accounting meters could be used to keep track of the responsible parties and to distribute charges appropriately (p. 2).

For higher-layer protocols, "there is a clear need to measure accounting variables and communicate them to various points along a packet's path" (p. 2). Accounting information could be communicated to a client in order to inform the user of charges incurred (p. 2). For old applications, unaware of the new accounting functions, accounting-aware gateways could be used for the application host in order to set the appropriate accounting records (p. 2). Since the tracking of accounting information can be assumed to take place in a hostile environment, appropriate measures will need to be taken to protect this information (p. 3). Network operators will need to plan for this and make any necessary changes (p. 3).

In order to be able to perform accounting at the level of IPng, "all that is required is the ability to carry one or two variable-size objects in every packet" (p. 3). Moreover, to protect the integrity of accounting information, encryption and digital signatures should be used (p. 3). Routers and accounting meters will need to have the appropriate encryption keys to be able to use this accounting information (p. 3). The fine details of accounting implementation can be standardized by the Accounting Working Group (p. 3).

Technical Code Summary

- Accounting support, direct
- Security support, direct

G. Electric Power Research Institute

The Electric Power Research Institute is a non-profit organization that manages technical research and development for the electric utility industry (Skelton, 1994).

Actor-Network, Participant Interests, and Proposed Solutions

The Internet Protocol is a critical element of the National Information Infrastructure (p. 1) and its redesign "provide[s] an ideal opportunity for creating the national uniform information transport superhighway" (p. 2). Although the electric utility industry is currently a heavy user of the Internet Protocol Suite, based on a survey of its needs, it has developed a long term strategy based on national and international standards (p. 1). Based on these standards and the OSI suite of protocols, the Electric Power Research Institute has developed the Utility Communications Architecture specification for communication and the Database Access Integrated Services specification for data exchange, both of which have been incorporated into the Industry Government Open Systems Specification (pp. 1-2). Both specifications have also been well received by the industry, its suppliers, as well as the natural gas and waterworks industries (p. 2). Consequently, in order to be successful, IPng proposals must be acceptable to both the private and public sector.

More specifically, IPng should be compatible with "upper layer industrial OSI applications" (p. 2). In addition, IPng should also include the following features. First, IPng should scale well and have "essentially an unlimited address space" (p. 2). Second, routing should be efficient and should include certain key features, such as "route aggregation, service selection, and low frequency host advertisements" (p. 3). Third, IPng should be able to interoperate with IPv4 and existing inter-domain routing protocols as well as be able to support proprietary protocols (pp. 2, 3). Fourth, the allocation of addresses should be based on considerations other than network topology (p. 3). Fifth, a method should exist for identifying network resources instead of network interfaces (p. 3). Dynamic addressing support should also be available to enable hosts to determine their network addresses and related network parameters (p. 3) Lastly, "[t]he solution must be available now using mature, internationally agreed standards and off-the-shelf implementations for hosts and routers" (p. 3).

Technical Code Summary

- Upper layer OSI application support
- Scalable addressing
- Address independent of network connection
- Network resource identification
- Route aggregation
- Service selection
- IPv4/IPng interoperability
- Support for other protocols
- Automatic configuration
- Available as soon as possible
- Based on OSI protocols

H. Cellular Digital Packet Data Consortium of Service Providers

The Cellular Digital Packet Data Consortium of Service Providers represents providers of digital packet data services for mobile devices using the Internet (Taylor, 1994).

Actor-Network, Participant Interests, and Proposed Solutions

Unlike location bound devices, mobile devices are not necessarily associated permanently with any particular geographical location. IPng should not inhibit the mobility of devices and should be connectionless. Moreover, since IPng provides an opportunity to redesign the Internet Protocol, "native support for mobility [should] be regarded as an explicit IPng requirement" (p. 1). Addresses should be independent of location and should specify an address, not a route through the Internet (p. 2).

A mobile "user" can be anything, "from a vending machine to a 'road warrior'" (p. 2). In the case of the latter, commercial users have security concerns regarding their communication (p. 2). These concerns are expected to be increasingly important in the future (p. 2). IPng should support "at least tens or hundreds of billions of addresses" worldwide (p. 2). Since commercial users demand security, IPng should have support for "peer entity authentication, data confidentiality, traffic flow confidentiality, data integrity and location confidentiality" (p. 2).

Since providers of digital packet data services do provide users with services, they charge users for the services rendered (pp. 1, 2). IPng should enable service providers to perform Internet Protocol level accounting as outlined in the Cellular Digital Packet Data accounting specification (p. 2).

In terms of the underlying infrastructure, mobile devices rely on both wide area wireless networks with limited bandwidth as well as low speed networks (p. 3). In order to optimize use of limited network bandwidth and low speed networks, IPng must allow for the optimal use of its underlying layer medium and support packet prioritization mechanisms (p. 3). In addition, IPng should provide service providers with a means of selecting routes that is compatible with existing voice communication practices (p. 2). Lastly, in order to remain compatible with existing IPv4 deployments, IPng must be able to "interoperate with [IPv4] for the foreseeable future" (p. 3).

Technical Code Summary

- Datagram based
- Address independent of network connection
- Scalable addressing
- Security support, direct
- Accounting support, direct
- Support for low bandwidth networks
- Resource reservation
- Policy routing
- IPv4/IPng interoperability

I. IPng Security

Steven Bellovin (1994) suggest that while IPng does not need to be more secure than IPv4 it cannot be less secure than IPv4.

Actor-Network, Participant Interests, and Proposed Solutions

IPng must be at least as secure as IPv4 (p. 1). As such, it must be compatible with existing firewall technology, which will continue to be utilized "to compensate for failings in software engineering and system administration" (p. 1). IPng will be less acceptable in the market if it makes firewalls more difficult to deploy (p. 1). In order to be compatible with existing firewalls, IPng must have a hierarchical address space, which firewalls rely upon for address-based filtering (p. 1). IPng must also make available the source and destination address, since firewalls often need to check these values (p. 2). Thirdly, IPng needs to enable firewalls access to transport-level (i.e. UDP, TCP) header information, since firewalls rely on this information to distinguish incoming and outgoing connections (p. 2).

If the transition from IPv4 to IPng will require the use of network-level packet translators, organizations that use firewalls will need to use their own translators, since they cannot rely on the centrally-located translators intended to serve the entire Internet (p. 2). Network-level translators will also need to be "simple, portable to many common platforms, and cheap," (p. 2) since "we do not want to impose too high a financial barrier for converts to IPng" (p. 2).

There are currently people who are experimenting with IP-level encryption and authentication. The experimentation and trend "will (and should) continue" (p. 2). To accommodate people who are experimenting with IP-level encryption and authentication, IPng should not make this activity harder "either intrinsically or by imposing a substantial perforance [sic] barrier" (p. 2). IPng must make possible encryption and authentication at various granularities, such as host to host, host to gateway, and gateway to gateway (p. 2). If hosts have support for both an IPv4 and an IPng TCP/IP stack, implementations should not allow hosts to bypass encryption or authentication by asking for a different address for the same host (p. 3). Likewise, if IPng has support for source-routing, cryptographic authentication should also exist, since source-routing will break IPv4 address-based authentication (p. 3).

Lastly, "[a] significant part of the world" wishes to do accounting, which "may be for billing, or it may simply be to accomodate [sic] quality-of-service requests" (p. 3). IPng should provide mechanisms for packet authentication and should allow billing systems to determine the relevant address fields as needed (p. 3).

Technical Code Summary

- Hierarchical addressing
- Source and destination addresses access
- Transport header access
- IPv4/IPng interoperability
- Security support
- Security support, direct
- Accounting support

J. Italian National Institute for Nuclear Physics Network Team

The Institute for Nuclear Physics Network Team manages the network of the Institute for Nuclear Physics in Italy (Ghiselli, Salomoni, & Vistoli, 1994).

Actor-Network, Participant Interests, and Proposed Solutions

IPv4 has three main problems: address exhaustion, flat address space, and inefficient, inflexible, routing with capacity shortcomings (p. 1). IPng should strive to correct each of these problems and not just some subset of them (p. 1). Although attempts should be made to fix these problems, user applications should not be affected (p. 2). In order for the changes to be transparent to user applications, the transport layer (i.e. TCP) should remain unchanged (p. 2). In addition, IPng should provide support for alternate transport layers and should include a "transport selector field" (p. 2).

Like IPv4, IPng should continue to provide a simple, connectionless datagram service (p. 2). The Open Systems Interconnect Connectionless Network Protocol could be either a solution or a good starting point for IPng, since "this solution has been profitable [sic] tested and it is already available on many systems" (p. 2). Any new features that are compatible with a datagram design should be included (p. 3). Source-routing, however, does not satisfy this requirement (p. 3). IPng should also not make any assumptions about its underlying media, which are quite variable and continue to be developed (p. 3). Network media should be allowed to develop unimpeded by IPng (p. 3).

The IPng routing protocol design is not fundamental, since routing protocols can always be changed or improved (p. 3). However, IPng routing protocols should include policy-based features, especially for the transition from IPv4 to IPng, which will be “very difficult or impossible to manage” (p. 3) without these features.

Based on experience with the DECNET network transition, every organization that needs to transition should be allowed to do so at its own pace. However, the need to support both an IPv4 and an IPng infrastructure indefinitely should be avoided. Although organizations should support systems with dual IPv4 and IPng functionality for a period of time, this support should only be provided during the transition from IPv4 to IPng (p. 3). Simultaneous support for IPv4 and IPng for a limited time is preferred over network translation mechanisms and tunnelling mechanisms to IPv4-only islands (p. 3). The former requires too much coordination and the latter will lead to poor performance (p. 3).

Technical Code Summary

- Scalable addressing
- Hierarchical addresses
- Transport layer unchanged
- Support for alternate transport layers
- Datagram based
- New capabilities relative to IPv4
- Heterogeneous subnetwork support
- Policy routing
- Transition autonomy
- Transition deadline

K. U.S. Navy Radio Frequency Tactical Systems Group

The U.S. Navy manages several projects that apply internetworking technology to radio frequency networks (Adamson, 1994).

Actor-Network, Participant Interests, and Proposed Solutions

U.S. Navy Internet applications typically “will include potentially very large numbers of local (intra-platform) distributed information and weapons systems and a smaller number of nodes requiring global connectivity” (p. 1). Within a military platform, there is a need to manage and control many distributed systems, such as radio communications equipment and weapons systems, for example (p. 3). Thus, IPng does not necessarily need to have support for many more globally-unique addresses, so much as support for the delimitation of globally-unique addresses from locally-unique addresses (p. 3).

The U.S. Navy maintains and supports many custom-built systems (p. 3). Some of these systems are being moved to TCP/IP technology (p. 3). Given this, IPng should be available for deployment as soon as possible to avoid having to perform two transitions, the first from custom technologies to IPv4 and the second from IPv4 to IPng (p. 3). However, the Department of Defense already maintains a large number of IPv4 systems that will need to be converted to IPng and "the issue of transition from IP to IPng remains significant" (p. 3).

Given the sensitive nature of military information, support for data confidentiality and authenticity "is of paramount importance" (p. 3). In addition, "[t]he need for network layer security mechanisms becomes more critical when the military utilizes commercial telecommunications systems or has tactical systems inter-connected with commercial internets" (p. 4). Although the U.S. Navy currently uses the Network Encryption Server, there exists "a desire for a more integrated, higher performance solution in the future" (p. 4).

Most tactical systems involve mobility of some sort, since members of tactical teams often need to join and leave particular radio frequency (RF) subnetworks (p. 4). In some cases, IPng will not need to have support for mobility, since the RF subnetwork itself will perform the necessary management of this mobility (p. 4). In other cases, however, major position changes may take place and "IPng must be able to support this situation" (p. 5). Moreover, mobility is not limited to individual nodes changing positions, but also includes the movement of entire subnetworks, as when "a Navy ship with multiple LANs on board, mov[es] through the domains of different RF networks" (p. 5).

The U.S. Navy makes extensive use of multimedia applications "from digital secure voice integrated with applications such as 'white boards' and position reporting for mission planning purposes to low-latency, high priority tactical data messages" (p. 5). Due to the quality of service requirements of these multimedia applications, IPng should provide support for resource reservation (p. 5). Given the limited capacity of RF networks, resource reservation mechanisms can be used to ensure that these quality of service requirements are satisfied (p. 5). However, based on experience and experimentation with RF networks, simple quality of service and source routing are inadequate for certain multimedia applications, such as real-time voice communication (p. 5). Instead, a more general technique such as the Resource Reservation Protocol's Flow Specification along with Flow Identification may be more appropriate for limited capacity RF networks (p. 6). Quality of service mechanisms and policy routing are still useful, however, to ensure that higher priority users are able to utilize the available capacity (p. 7). IPng should support the assignment of priority "to users or even individual datagrams" (p. 7). In a similar vein, IPng headers should not grow "overly large", since there is a great trade-off in performance between compact and large packet headers (pp. 3, 8). Even though the capacities of future RF networks may exceed current capacities, "there is always a tactical advantage in utilizing your resources more efficiently" (p. 8).

In addition, since "[m]any of the tactical RF communication media are broadcast by nature" (p. 7) and are used to "distribute critical data to a large number of participants" (p. 7), IPng should support multicasting. Limited capacity RF networks and "the physics of potential electronic counter measures (ECM) dictate that this information be distributed efficiently" (p. 7).

Like IPv4, IPng should provide a datagram service, since state is maintained in packet headers and "provides an inherent amount of survivability essential to the dynamics of the tactical communication environment" (p. 7). IPng should also be able to coexist with other architectures and should support all forms of underlying communication media, since the military owns and operates a wide variety of media and technologies that need to interact with each other (pp. 7-8).

Technical Code Summary

- Scalable locally-unique addresses
- Available as soon as possible
- Transition autonomy
- Security support
- Security support, direct
- Mobile hosts
- Mobile networks
- Mobile internetworks
- Resource reservation
- Service selection
- Policy routing
- Compact headers
- Multicasting
- Datagram based
- Support for low bandwidth networks
- Link sharing with other protocols
- Heterogeneous subnetwork support

L. Corporations #1

"As more and more corporations are using TCP/IP for their mission- critical applications, they are bringing additional requirements … the satisfaction of which would make TCP/IP even more appealing to businesses" (Britton & Tavs, 1994).

Actor-Network, Participant Interests, and Proposed Solutions

Since the Internet is very large today, all of its users cannot be expected to transition at the same time (p. 2). In addition, users will not switch to IPng unless it is able to provide new services that are unavailable in IPv4 (p. 2). Likewise, users will not switch, or may not be able to switch, if IPng requires significantly more resources (i.e. memory, storage, etc.) than IPv4 (pp. 2-3). As a result, IPng needs to be able to coexist with IPv4 and should provide transition mechanisms that enable Internet users to switch at their own pace (pp. 1-3).

Corporations rely on TCP/IP for applications that are critical to their business activities (pp. 1, 2). Similarly, corporations use budgets and service level agreements to plan for the costs of technology and expected network services (p. 3). Consequently, IPng "should allow control of the cost of networking, a major concern for corporations" (p. 3). In addition, since teleprocessing lines are expensive, corporations expect that IPng will allow maximum utilization of these lines, "but without requiring excessive processing power to achieve the high utilization" (p. 3). Support for functions such as class of service and traffic priority may help corporations better control the use of teleprocessing lines (p. 3). More generally, to discourage waste of expensive resources, including bandwidth, IPng should allow corporations to account for network use (p. 3). Corporations also expect to be able to receive varying grades of network service for varying costs (p. 4).

Until public networks can guarantee security, confidentiality, and integrity that is comparable to private networks, corporations will continue to rely on their private networks (p. 4). IPng may not necessarily be the solution to the problems of security experienced by corporations. Other security technologies may suffice (p. 4). However, IPng should not interfere with firewalls, which many corporations rely on for security (p. 4). In addition, since corporations "do not want to carry for free the transit traffic of other enterprises, and they may not want their sensitive data to flow through links controlled by certain other enterprises" (pp. 6-7), IPng should include support for policy routing that would enable corporations to select their traffic flows (pp. 6-7).

Corporations often use several protocol suites (pp. 2, 5). Corporate decentralization, technical factors, workforce skills, and diverse installed bases contribute to the use of multiple protocol suites (p. 2). As a result, "IPng should be flexible enough to accommodate a variety of technical approaches to achieving heterogeneity" (p. 5). It should accommodate multiprotocol transport networking, tunnelling, and link sharing with other protocols (p. 5).

Corporations are also interested in being able to use fast packet switching networks and their corresponding multimedia services (p. 5). Thus, IPng should provide support for constant-bit-rate and variable-bit-rate services over the same link and should include mechanisms to "discourage inappropriate reservation of resources" (p. 6).

Wireless technologies provide new opportunities and challenges to TCP/IP applications (p. 6). IPng should "deal well with the characteristics (delay, error rates4 [sic], etc.) peculiar to wireless" (p. 6). Currently, IP address and proxy servers are used to overcome the limitations placed by IPv4 on mobile hosts (p. 6). In addition, moving non-wireless hosts is restrictive and complicated under IPv4 (p. 6). Consequently, "IPng needs an addressing model more flexible than subnetting, both to reduce the administrative burden and to facilitate roaming users" (p. 6).

Technical Code Summary

- IPv4/IPng interoperability
- Transition autonomy
- Resource reservation
- Service selection
- Predictable service for predictable cost
- Accounting support
- Security support
- Policy routing
- Support for other protocols
- Link sharing with other protocols
- Support for multiprotocol transport networking
- High performance subnetwork support
- Mobile hosts
- Scalable addressing

M. U.S. Navy High Performance Network Group

The U.S. Navy High Performance Network Group researches network architectures for applications aboard U.S. Navy platforms such as aircrafts, ships, and submarines (Green, Irey, Marlow, & O'Donoghue, 1994).

Actor-Network, Participant Interests, and Proposed Solutions

The U.S. Navy relies on mission critical network applications in a distributed combat environment (p. 2). Although the U.S. Navy makes use of current Internet and ISO network protocols, they are inadequate for the Navy's mission critical applications (p. 2). In general, the U.S. Navy's distributed combat system environment includes "a collection of workstations involved in many varied applications involving multiple sources and types of information" (p. 2). In addition, the internetworking environment includes many heterogeneous systems that make use of various subnetworks (p. 2). In these environments, a wide variety of network endpoints, such as devices and end users need to be supported (p. 2).

Although the current Internet network protocols lack certain important features, IPng should, as a minimum, support the existing functionality of IPv4 (p. 4). That is, IPng should have mechanisms that provide connectionless transfer, support multiple subnetwork media, support hosts that utilize multiple subnetwork media, and support network management functions (p. 4). In addition, IPng should be able to coexist with other network protocols (p. 4). As with IPv4, IPng should also have the support of the commercial/industrial community. It should function correctly and should demonstrate multi-vendor interoperability (p. 4). Given that the U.S. Navy maintains both proprietary networks as well as IPv4 networks, IPng should provide mechanisms that enable transition from these technologies (p. 5).

A single U.S. Navy platform will typically include anywhere from 100 to 100,000 addressable entities and many of these platforms will be connected to global networks (p. 5). Consequently, IPng must be able to support many globally-unique addresses. In addition, IPng should provide support for logical addressing that is independent of network topology (p. 5). Multicast group addressing should also be possible: "At a minimum 2**16 globally unique multicast groups must be distinguishable per platform" (pp. 5-6). Given the nature of many platform activities, shared data is needed by multiple hosts and the ability to efficiently manage and perform multicast transfers is highly desirable (p. 7). Current IPv4 multicast implementations are inadequate (p. 7). Platforms also include mobile hosts, such as mobile terminals, that often change their location (p. 6). Likewise, the platform itself may be mobile, as in the case of an aircraft carrier (p. 5). As a result, IPng needs to be able to support mobile hosts, networks, and internetworks.

Given the type of information used, which includes voice, audio, video, image data, text, and sensor data, IPng must provide support for these various information flows and their performance requirements (pp. 3, 6). In addition, IPng should provide mechanisms that enable the management of the various traffic flows associated with different types of information and services (p. 6). That is, in addition to the requirements of various information services, it should be possible to prioritize traffic flows to ensure that high priority information arrives where it is needed as fast as possible (p. 6).

The U.S. Navy utilizes very high bandwidth subnetwork technology (p. 7). Given the mission critical nature of the applications that rely on the subnetwork, IPng should include support for mechanisms that enable fast route reconfigurations when subnetwork problems arise (p. 7). Specifically, subnetwork fault detection/reconfiguration should take "less than 1 second" (p. 7). In general, IPng support for "error detection, error reporting, traffic analysis, and status reporting" (p. 8) is desirable.

In terms of security, the U.S. Navy's mission critical applications require both data security and integrity and IPng should support a number of security functions, including "rule based access control, labeling, authentication, audit, connection oriented and connectionless confidentiality, selective routing, traffic flow confidentiality, connection oriented and connectionless integrity, denial of service protection, continuity of operations, and precedence/preemption" (p. 8). IPng should include support for IPSEC as well as the ability to add future security functions (pp. 8-9).

Lastly, since many of the U.S. Navy's mission critical applications need to be time synchronized, IPng should at a minimum provide support for higher layer time synchronization algorithms (p. 8). Specifically, "it is desirable that a time-of-day synchronization capability be supported of at least an accuracy of one microsecond among all hosts in a platform" (p. 8).

Technical Code Summary

- Datagram based
- Heterogeneous subnetwork support
- High performance subnetwork support
- Link sharing with other protocols
- Multi-homed host support
- Support for network management
- Commercially viable
- Transition management tools
- Mobile hosts
- Mobile networks
- Mobile internetworks
- Resource reservation
- Fast route reconfiguration
- Error reporting
- Policy routing
- Security support

- Scalable addressing
- Address independent of network connection
- Multicasting
- Security support, direct
- Support for time synchronization

N. Asynchronous Transfer Mode

Typical link technologies such as Ethernet simply multiplex traffic on a single channel with a single performance and quality of service characteristic. In contrast, a "sophisticated" link technology such as Asynchronous Transfer Mode (ATM) enables the establishment of several virtual channels, each with its own performance and quality of service characteristics (Brazdziunas, 1994).

Actor-Network, Participant Interests, and Proposed Solutions

IPng needs to support both existing and "sophisticated" link technologies such as ATM, since IPng packets will traverse both types of subnetworks (p. 1). More specifically, "[t]hough initial support for IPng over ATM subnetworks will not facilitate a virtual circuit per application, the hooks to provide such a mapping should be in place" (p. 1). It should be possible to attain the following information from protocols above the link layer (including IPng): "source and destination(s) addresses, connection quality of service parameters, connection state, and ATM virtual circuit identifier" (p. 1). With the exception of the ATM virtual circuit identifier, which should be derivable from an IPng packet, the required information may be derivable from currently proposed resource reservation protocols (p. 1).

Technical Code Summary

- ATM support

O. Multiprotocol Interoperability In IPng

Clark, Ammar and Calvert (1994) argue that IPng should incorporate features that enable multiprotocol interoperability.

Actor-Network, Participant Interests, and Proposed Solutions

In addition to the Internet protocols, the Internet includes a variety of other protocols that either coexist alongside the Internet protocols or that are encapsulated as data and transmitted using the Internet protocols (p. 1). Multiprotocol networking is not simply a transition issue, since current protocols will continue to be used in addition to IPng and new protocols will continue to be developed (pp. 1-2). Consequently, IPng should support functionality that improves its ability to operate in a multiprotocol environment (p. 1).

Since multiprotocol techniques such as tunnelling and translation/conversion have serious limitations, IPng should include support for explicit protocol determination (p. 5). That is, the Domain Name System could be modified so that in addition to returning address information for a particular host, protocol configuration information could also be provided (pp. 5, 9). Based on its own protocol support and communication needs, the host could then use this protocol information to select the appropriate protocol for the communication. Alternatively, a trial-and-error method could be used to try to discover which protocols were supported by the destination host (p. 5).

Given that most protocols use a unique addressing scheme, IPng should be flexible enough to accommodate as many of these address formats as possible with as little additional address conversion as necessary (p. 6). The Domain Name System should also be modified to support address formats for protocols other than IP (p. 9). Also, since there will be many protocols running on top of IPng, it should not be necessary to use a transport layer to provide multiplexing (p. 7).

In addition, to support a wide variety of users with a wide variety of needs, IPng should have flexible header options that enable the specification of options as they are needed or as they are developed in the future (p. 7). As part of this support, IPng should specify how a system should handle unsupported options (p. 7). IPng should also include "a powerful error reporting mechanism," since if tunnelling or conversion are involved there are "many different places the transmission can fail and determining what went wrong will be a challenge" (p. 8).

Like other protocols, IPng will need to run on diverse link and network layers and, consequently, should be designed to support different underlying link layers (pp. 8, 9). Users should also have the option to specify which of the underlying link layer services should be used (p. 9). To work well within the limits various Maximum Transmission Units, IPng will need mechanisms to optimize transmission sizes, since both tunnelling and conversion may lead to larger packets and additional fragmentation (p. 8).

Technical Code Summary

- Support for other protocols
- Link sharing with other protocols
- Heterogeneous subnetwork support
- Multi-homed host support
- Extensible header
- Error reporting
- Optimal transmission size determination

P. The Cable Television Industry

The cable television industry designs and deploys integrated broadband networks to support new and existing services and applications (Vecchi, 1994).

Actor-Network, Participant Interests, and Proposed Solutions

Cable operators are deploying integrated broadband networks and plan to extend cable services to encompass new applications in addition to traditional television offerings (pp. 1, 2). That is, "the architecture of future cable networks maps directly to the way networked computing has developed" (p. 4) in that "[g]eneral purpose hosts (i.e., the set-top boxes) are interconnected through a broadband network to other hosts and to servers" (p. 4). Given that compressed digital video is "the way to deliver future video programs" (p. 4) such as interactive video, video on demand, multiparty remote games, computer supported collaborative work, home shopping, customized advertisement, and multimedia information services, "one can be guaranteed that gigabit regional networks will be put in place at an accelerated pace" (p. 4).

Based on the number of households worldwide that can be expected to subscribe to cable services, IPng will need to support at least 10 to the power of 10 globally unique addresses (p. 5). In addition, since cable operators and their partners will be making decisions with respect to investments and deployments in the near future (12 to 24 months), the features and capabilities of IPng should be decided upon within this timeframe (p. 6). Transition mechanisms should be in place that support the transition from IPv4 to IPng in such a way that existing IPv4 applications can be supported (pp. 6, 7). Like IPv4, IPng should have support for a variety of underlying communication media (p. 12). In order to ensure network robustness during the transition, "[t]he design, implementation and testing process will have to be managed very carefully" (p. 12).

Although IPng should be backwards-compatible with IPv4, it should provide new capabilities, the most important of which is support for high quality video and audio streams with real-time constraints such as end-to-end delay bounds (p. 6). Delivery of audio and video services requires that IPng support broadcasting and selective multicasting (p. 6). Likewise, resource reservation mechanisms will need to be in place to ensure quality of service (pp. 7-8). Given that there will be "multiple network operators and information providers competing for customers and network traffic" (p. 10), IPng will need to support source routing so that users "will be able to select specific networks that provide performance, feature or cost advantages" (p. 10). For cable operators, source routing "would enable the offering of targeted bundled services that will cater to specific users and achieve some degree of customer lock-in" (p. 10). IPng should also provide support for accounting either at the network layer or above, depending on the performance overhead required (p. 11), since "future broadband networks will be commercially motivated, and measurement of resource usage by the various users will be required" (p. 11).

In addition, since both users and cable operators will expect a high level of security, IPng should have security mechanisms that protect data traffic. That is, "[t]he possibility of illicitly monitoring traffic patterns by looking at the headers in IPng packets, for instance, could be disturbing to most users that subscribe to new information and entertainment services" (p. 7). Given that cable operators will need to configure, administer, and operate increasingly complex applications and networks, IPng should also either support, or not get in the way of, capabilities that enable operators to "dynamically reconfigure the connectivity among end points or the location of specific processes" (p. 8). Lastly, IPng should have support for mobile devices, which are "being introduced to the market at an accelerated pace" (p. 8).

Technical Code Summary

- Scalable addressing
- Available as soon as possible
- IPv4/IPng interoperability
- Heterogeneous subnetwork support
- Multicasting
- Resource reservation
- Source routing
- Accounting support
- Security support, direct
- Support for network management
- Mobile hosts

Q. Corporations #2

Fleischman (1994) outlines the official position of the Boeing Company with respect to IPng. In addition, "[a]n assumption of this paper is that other Fortune 100 companies which have non-computing-related products and services will tend to have a viewpoint about IPng which is similar to the one presented by this paper" (p. 1).

Actor-Network, Participant Interests, and Proposed Solutions

Fortune 100 corporations, understood as end users of networking technology, have invested heavily into TCP/IP (p. 2). Although the IETF community views IPng in positive terms, most corporations are not enthusiastic about IPng (p. 2). Unlike the IETF community, which closely associates TCP/IP with the Internet, end users such as corporations see the relationship between TCP/IP and the Internet as "tenuous at best" (p. 3). This is due to the fact that corporations have made substantial investments into TCP/IP for internal networks that do not necessarily have any Internet connectivity at all (p. 3). In addition, many corporations have also invested heavily into external communication infrastructures to support their business functions, since until recently the Internet was excluded from bilateral agreements and international tariffs required for international commerce (p. 3). As a result, "end users today are not (in the general case) dependent upon the Internet to support their business processes" (p. 3), since external communication infrastructures fulfill this need. However, corporations that have invested into TCP/IP for internal networking "may perhaps view TCP/IP technology to be critically important" (p. 3).

Unlike the IETF, which "generally has a world view which presupposes that data access should be unrestricted and widely available" (p. 4), corporations view security islands as a necessary mechanism to protect their data and computer resources both from external parties, who should have either no access or limited access to data, and internal parties such as employees, who have varying data and access needs and permissions (p. 4). Consequently, corporations need to be able to define their own security islands by means of firewalls and other technologies (p. 5).

Given that many corporations limit the number of hosts that are directly accessible from the Internet, "the state of the world-wide (IPv4) Internet address space only directly impacts a tiny percentage of the currently deployed TCP/IP hosts within a large corporation" (p. 5). That is, the problem of IPv4 address depletion is not important for many corporations, since their internal networks are unaffected. In addition, since most corporations have many more hosts than routers, a transition from IPv4 to IPng that only involves the updating of routers "is generally much less disruptive to a corporation than an approach which demands changes to both routers and hosts" (p. 5).

Likewise, if IPng introduces an entirely new set of protocols it will likely not be adopted by corporations, since it will increase corporations' overall protocol complexity (p. 6). Most corporations already support many different protocol families and strive for interoperability between protocol families to increase the efficiency of available computer resources (p. 6). The introduction of a new protocol family also “potentially implies unique training, administrative, support, and infrastructure requirements” (p. 6). Unless there are “significant business benefits” (p. 7) to justify the deployment of IPng, “economics will oppose such a deployment” (p. 7).

Corporations use computer applications, supported by networking technology, to automate business processes and increase efficiency. Unless IPng can support new or significantly improved applications in this regard, corporations are not likely to deploy it (p. 6). In the future, if the Internet requires IPng usage or if there are eventually many remote peers that solely use IPng, corporations may only partially deploy IPng in their gateway systems (p. 7). Similarly, if senior management directs IPng adoption or if IPng enables new or significantly improved applications, corporations may only deploy IPng on those internal networks where it is needed or mandated (p. 8).

Despite current usage and deployment of IPv4, there are a number of areas that could be improved upon by IPng in order to increase the likelihood of adoption by corporations (pp. 2, 9). First, given that TCP/IP is typically deployed over broadcast media such as Ethernet, IPng should have improved support for security and should include encryption mechanisms (p. 8). Secondly, given the existing difficulties, problems, and errors with IPv4 configuration, IPng should support auto-configuration and plug-and-play TCP/IP networks that “auto-configure, auto-address, and auto-register” (p. 9). Thirdly, since corporations deploy multiple protocol families, tools are needed to integrate these protocol families and to aid with protocol transitions (p. 9). Fourthly, corporations are still not sure whether TCP/IP “can be universally used for international commerce” (p. 9). Lastly, corporations are interested in improving existing applications and technologies. An IPng that is able to support mobile hosts, multimedia and real-time applications, high-bandwidth applications, very low-bandwidth applications, and transaction-based applications may be more readily adapted by corporations (p. 9).

Although several IPv4 problem areas present opportunities for IPng, corporations are ultimately most interested in support for new and improved applications and the ability to integrate TCP/IP with existing and emerging international standards that include the OSI protocols (p. 11). Integration is important both to improve interoperability with existing installations that make use of international standards as well as to reduce the costs involved with the management of diverse protocol families (pp. 11-12).

Technical Code Summary

- Transition autonomy
- Scalable locally-unique addressing
- Security support
- Transition via routers only
- New capabilities relative to IPv4
- Security support, direct
- Security improvements relative to IPv4
- Automatic configuration
- Transition management tools
- Support for international commerce
- Mobile hosts
- Resource reservation
- Support for high bandwidth networks
- Support for low bandwidth networks
- Datagram based
- Integration with international standards
- Based on OSI protocols
- Support for other protocols

R. IETF Mobile IP Working Group

The IETF Mobile IP Working Group has developed criteria related to mobility for consideration in the design of IPng (Simpson, 1994).

Actor-Network, Participant Interests, and Proposed Solutions

Currently, the Internet Protocol assumes that a node's location, or point of attachment, remains fixed (p. 2). As a consequence, if a node does change location while keeping its IP address, data cannot be properly routed to it, since its IP address will not reflect its new point of attachment (p. 2). In order to overcome this limitation, several changes need to be made to accommodate mobility.

To begin with, each mobile node should have at least one home address that identifies it uniquely to all other nodes (p. 2). The home address could have ownership information associated with it that could then be used to "facilitate 'inverse' lookup in the Domain Name Service, and other future services" (p. 2). Other nodes could use the ownership information to determine the current topological location of the mobile node (p. 2). The ownership information could also be used for accounting purposes (p. 2). However, the actual home address need not include any topological information at all, and if it does, it should not in any way hinder the mobility of the node (p. 3). Likewise, if manufacturing information is associated with the home address, it should not in any way hinder node mobility (p. 3).

It is assumed that the number of mobile nodes will be limited by the population of users, which is "not expected to exceed 8e9" (p. 3) during the lifetime of IPng. Given that each user is expected "to have a maximum combined personal mobile and home network(s) on the order of 4e3 nodes" (p. 3), only 46 bits will be required to number all mobile and home nodes (p. 3). In addition, "[s]ince the typical user would be unlikely to be aware of or willing and able to maintain 4e3 nodes, the assignment of Home-Addresses must be automatically configurable" (p. 3).

Given that the network links of mobile nodes are "likely to be bandwidth limited" (p. 6), mobility "must not be dependent on bandwidth improvements" (p. 6). Similarly, mobility must not depend on improvements in hardware, since typical and existing mobile nodes rely on lower speed processors that are "not readily upgradable" (p. 8). These processing limitations require that IPng header fields, which are frequently examined by mobile hosts, are in a fixed location, are fixed in number, and do not make use of lengths and offsets (pp. 8, 9). Likewise, given that mobile node processors are typically little-endian, whereas network protocols are big-endian, field processing "must primarily use simple equality tests, rather than variable shifts and prefix matches" (p. 9). Lastly, IPng headers should either be no larger or only slightly larger than IPv4 headers in order to "maintain an acceptable typing response of 200 milliseconds round trip time" (p. 7).

Technical Code Summary

- Address independent of network
- Support for low bandwidth networks

connection
- User information associated with addresses
- Mobile hosts
- Scalable addressing
- Automatic configuration
- Fixed size headers
- Fixed location header fields
- Simple equality tests for header processing
- Compact headers

S. Multiple Addresses Per Host

Bellovin (1994) argues that address space calculations incorrectly assume that each host will only have one IP address. The calculations should allow for the possibility that each host could have multiple IP addresses.

Actor-Network, Participant Interests, and Proposed Solutions

There are several areas of application where multiple IP addresses per host may be useful or may provide a solution to existing problems (p. 1). First, the existing practice of associating a host name with a service, such as ftp.research.att.com, for example, could be extended to the association of an IP address with a service (p. 1). In this way, service providers could offer multiple views of a given service on a single host "without asking the user community to learn new protocols or custom port numbers" (p. 2). If the services have different access restrictions, a firewall could be used to control access to one view while allowing access to the other (p. 2). Although a similar effect could be achieved by modifying the Domain Name Service so that it would return port numbers in addition to IP addresses, such a change would require that all client applications be modified in order to process the new Domain Name Service information (p. 2). Likewise, if a service provider does not wish to rely on the Domain Name Service at all, associating IP addresses with services would allow service providers to easily move services from one host to another without affecting users, who would continue to use the same IP address to reach the service (p. 2).

A second area of application involves accounting and billing, since "[f]or better or worse, some parts of the Internet are moving towards usage-sensitive charging" (p. 3). Marketeers, building on telephone practices and models, are devising various charging schemes to ensure that all packets are accounted for (p. 3). Regardless of the charging scheme, “it is vital that the caller and recipient know in advance who will pay” (p. 3). Although new IP options could be devised, this would entail modifying all existing client applications so that they could work with and process the new options (p. 3). Likewise, messages at connection time that inform users of the charge rate are inadequate, since “[m]any interactions do not provide a hook for user interaction” and those that do are prone to scams, which are likely to happen given the history of telephone networks (p. 3). Instead, the solution is to “encode the charge algorithm in the destination address” (p. 3), offering differently charged services via different IP addresses.

Thirdly, it may be useful to associate a unique IP address with each Internet user, both to support accounting and billing as well as for other applications (p. 4). That is, “different classes of users could have different forms of addresses” (p. 4) such that users “with hard-money accounts might have some bits set in the address that would allow for access to costly services” (p. 4). Existing technology, and in particular border routers, could support these distinctions and bill each user appropriately (p. 4).

Lastly, support for multiple addresses per host would enable system administrators to setup load balancing schemes that dynamically move IP addresses from host to host based on service traffic and available resources (p. 4). Similarly, in situations where a host needs to exist on two or more logical networks, multiple IP addresses per host would enable this parallel existence without the need for traffic to cross routers (pp. 4-5).

To support potential multiple addresses per host application areas, “we [should] allow for 2^6, and perhaps as many as 2^8, extra addresses per host” (p. 5).

Technical Code Summary

- Multiple addresses per host
- Scalable addressing

T. Transport and Network Layer Independence

The existing transport and network layers are not protocol independent. Both TCP and UDP need to be redesigned in such a way that it is possible to use network protocols other than IP (Carlson & Ficarella, 1994).

Actor-Network, Participant Interests, and Proposed Solutions

During its development, TCP was strongly linked to IP (p. 2). Although this linking enables efficient Internet operation, it simultaneously makes it difficult to modify one protocol without modifying the other. This observation is exemplified by IPng proposals, each of which requires some changes at the transport layer in order to support features at the network layer (p. 3). The existing TCP/IP violates the OSI layer model, which "has now, in concept, become the fundamental building block of computer networks" (p. 3) and which requires that each layer be independent of the adjacent layers, such that it should be possible to change one layer without affecting the others (p. 7). That is, it should be possible to change each layer's implementation as long as it continues to satisfy the interface specification between layers (p. 7). The OSI layer model is well suited to present day computer and communication systems that "have become increasingly heterogeneous in both their software and hardware complexity" (p. 3).

IPng must replace IPv4 due to the inherent limitations of the IPv4 address space, which is both flat and limited in terms of available addresses (p. 4). A large scale transition from IPv4 to IPng will be difficult to accomplish for a number of reasons. First, there will be a need to "rewrite highly utilized and entrenched user applications, such as NFS, RPC, etc." (p. 5), since they internally embed specific IP addresses (p. 6). Second, new hardware and software will need to be purchased to support the new protocols (p. 5). Thirdly, and in addition to this capital reinvestment, the user community will need to be retrained in order to make use of the new protocols and systems (p. 5).

By replacing the current TCP layer with a next generation TCP layer (TCPng) that supports a unique transport layer address that is independent of the network layer, the problems associated with the transition from IPv4 to IPng will be minimized, since "[t]he net affect is that while administrators will necessarily be trained in the operations and details of this new TCP, the much larger operator and end user community will experience no perceptible change in service and network usage" (p. 5). That is, systems that make use of TCPng will still be able to connect to systems that support the old TCP (p. 6). Likewise, system administrators will be able to transition to IPng at their own pace, since systems with IPng, regardless of whether IPng is backward compatible with IPv4, will be able to make use of TCPng to communicate with hosts that still rely on IPv4 (pp. 6-7). Lastly, "[a]s nodes upgrade to TCPng, new features will be enabled allowing TCP to communicate effectively over high bandwidth*delay [sic] network links" (p. 6).

Although the IAB promotes a single Internet protocol suite, the addition of a transport layer address to TCP, which will support a multiprotocol Internet, "does not invalidate the IAB's stated goals" (p. 7), but "merely brings the implementation into compliance with standard networking practices" (p. 7).

Technical Code Summary

- Transport layer address independence

U. IETF Nimrod Working Group

The IETF Nimrod Working Group designs and develops the Nimrod routing and addressing architecture for the next generation Internet (Chiappa, 1994).

Actor-Network, Participant Interests, and Proposed Solutions

Current-day routing technologies "do not seem to be completely suited to routing in the next-generation Internet" (p. 2). As routing technology evolves to meet the demands of the next generation Internet, "the underlying fundamental laws and principles of routing will almost inevitably drive the design” (p. 2) towards a design that resembles the Nimrod routing and addressing architecture (p. 2).

The general design philosophy of Nimrod is “‘maximize the lifetime (and flexibility) of the architecture’” (p. 2). As a result, the design should be simple and should avoid optimizations that may eventually lead to unforeseen inflexibility (p. 2). Functions should be handled by independent mechanisms and packet fields should be individually specified (p. 2). In addition, field lengths should be specified with future expansion in mind (p. 3).

Aside from general design philosophy, Nimrod has three planned modes for forwarding packets: flows, datagram, and source routing (p. 3). IPng packets will need to include packet fields necessary to support these planned modes. Specifically, IPng should include the following: globally unique, topologically insensitive identifiers for network endpoints or multicast groups, flow identification, a looping packet detector to determine if a packet is “stuck” in the network, optional source and destination locators that are topologically sensitive, an optional pointer into these locators, an optional source route with an optional pointer into this source route, header length, support for authentication, and a version number to distinguish different Nimrod versions coexisting in the same system (pp. 4-8).

In addition to these specific packet requirements, flows should be explicitly supported by IPng in a unified manner (p. 17). It should also be possible to determine whether a given resource allocation request can be satisfied (p. 17).

Nimrod also requires that IPng support multiple interfaces per endpoint, multiple globally unique identifiers per endpoint, and multiple addresses per interface (p. 17). In addition, IPng should distinguish between multicast groups and multicast flows (p. 17).

Technical Code Summary

- Algorithm and data structure independence
- Header length
- Authentication

- Extensible header fields
- Unique names
- Flow identification
- Multicast
- Looping packet detector
- Hierarchical addresses
- Version identification
- Resource reservation
- Multi-homed
- Multiple unique names
- Multiple addresses

The IPng Area and RFC 1726: Translating Everyone

Following the publication of white papers solicited via RFC 1550, Partridge and Kastenholz (1994) wrote *Technical Criteria for Choosing IP The Next Generation (IPng)*, which was published as RFC 1726 and formed the basis for evaluating IPng proposals:

> These white papers, a Next Generation Requirements (ngreq) BOF (chaired by Jon Crowcroft and Frank Kastenholz) held during the March 1994 Seattle IETF meeting, discussions within the IPng Area Directorate and considerable discussion on the big-internet mailing list were all used by Frank Kastenholz and Craig Partridge in revising their earlier criteria draft [Kasten92] to produce "Technical Criteria for Choosing IP The Next Generation (IPng)." [Kasten94] This document is the "clear and concise set of technical requirements and decision criteria for IPng" called for in the charge from the IESG Chair. (Bradner & Mankin, 1995, p. 9)

In the above passage, the term "BOF" refers to an IETF "birds of a feather" session, which is used to denote an informal meeting on a particular topic. The BOF, along with the white papers solicited by RFC 1550, IPng Area Directorate discussions, and mailing list discussions were summarized by Partridge and Kastenholz such that RFC 1726 may be understood as translating the participant interests of all of the relevant social groups involved with IPng. That is, in addition to the participant interests of the IPng Area and the 21 relevant social groups that responded to RFC 1550, RFC 1726 also includes the translation of participant interests that were articulated at the Next Generation Requirements BOF and via the IETF mailing list. The document translates these participant interests into specific IPng criteria that IPng proposals need to satisfy. In this sense, RFC 1726 is very similar in form to the white papers solicited via RFC 1550 in that it traces an actor-network that requires certain attributes and actions from IPng.

Consequently, the following description of RFC 1726 follows the conventions used previously to describe the relevant social groups and their participant interests.

Actor-Network, Participant Interests, and Proposed Solutions

The primary purpose of the Internet is to provide an "IP Connectivity Service to all who wish it" (p. 4). Many communications functions are more appropriately performed at other protocol layers and the network layer should be architecturally simple (p. 4). When IPng does not provide a particular function, "it should not get in the way of the other elements of the protocol stack which may well wish to perform the function" (p. 4). In addition, since the Internet is decentralized and composed of many individual network constituents, IPng should continue to allow each constituent to "tailor their own networks, environments, and policies to suit their own needs" (p. 5).

Since the user community, consisting of vendors, service providers, and users, cannot be expected to undergo frequent and major protocol changes, IPng should be designed to accommodate the Internet of the future (pp. 5, 21). In order to provide incentive for the user community to switch from IPv4 to IPng, IPng should include support for new features and should include mechanisms that enable the Internet "to prosper and to grow to serve new applications and user needs" (p. 5). These mechanisms should include algorithm independence, data structure independence, extensible packet headers, and support for new packet types (p. 22). As IPng evolves, it must be possible for different versions to coexist on the same network (p. 21).

IPng must overcome the address and routing limitations of IPv4 (p. 7). To the extent that in the future each home will be a network, "there will be the need to support at least 10**12 networks, with the possibility of supporting up to 10**15 end-nodes" (p. 7). Moreover, network constituents should be able to structure their internal addressing and topologies in any manner they see fit (p. 8). Since the Internet is currently not structured as a tree and cannot be expected to be structured in this way in the future, IPng addressing and routing should also support multiply connected networks (p. 8).

IPng must be able to achieve data transfer rates comparable to IPv4 using similar levels of host resources (p. 9). IPng must also be able to support high speed network media and needs to be able to guarantee performance that is commensurate with network media and technology advances (p. 9). Improved header processing performance may be achieved by optimizing header field sizes for future processors and word widths (p. 10). In addition, IPng, like IPv4, should be robust despite occasional failures of hosts, links, and routers (p. 10). It should "perform gracefully in response to willful management and configuration mistakes" (p. 10) and should be able to handle malformed packets and misinformation (p. 10). The effects of failures should be localized as much as possible (p. 11).

Given the size, complexity, and distributed administration of the Internet, a "flag-day" transition is not possible. Instead, "[a] smooth, orderly, transition from IPv4 to IPng is needed" (p. 12). IPv4 hosts will continue to exist on the Internet long after IPng becomes the dominant protocol and IPng needs to be able to coexist and communicate with these hosts (p. 12). Likewise, IPng is expected to evolve over time and it should be possible for different versions of the protocol to coexist on the same network (p. 12).

In addition to supporting high speed network media, IPng should also be able to support a wide variety of subnetwork technology, including limited bandwidth media (p. 13). In this sense, IPng should function like IPv4: "The joy of IP is that it works over just about anything" (p. 13). In particular, IPng should include support for ATM, which "seems to be set to become a major network media technology" (p. 14). Since packets will likely traverse multiple media to reach their destination, IPng must be able to adapt to the variable transmission size constraints of these media (p. 14). Likewise, given that many network media support more than one protocol suite, IPng "must be able to peacefully co-exist on such media with other protocols" (p. 14).

As is the case with IPv4, IPng should provide a basic datagram service, since this feature is "extremely valuable" (p. 15) and required by certain applications (p. 15). By enabling packets to be dropped by routers and enabling routers to go offline without any warning, datagram service contributes to robustness and "vastly eases administration and maintenance of the network" (p. 15). In addition, it "vastly simplifies the design and implementation of software" (p. 15). IPng's datagram support should provide a minimal level of service that is similar to that provided by IPv4 (p. 15).

Since "[p]eople complain that IP is hard to manage" (p. 16), IPng must support automatic configuration of hosts and routers (p. 16). Although automatic configuration for large networks is complex and is still the focus of research, the requirement that IPng support automatic configuration is applicable especially to small and medium sized networks where user expertise and knowledge is limited (p. 16). In addition, it should be easy to allocate addresses and easy to change any of network topology, network provider, host parameters, or router parameters (p. 16). IPng should also not prohibit users from manually configuring their systems and should not obstruct other network configuration technologies from functioning (p. 17).

Given that the Internet operates in a hostile environment that includes crackers and eavesdroppers, IPng needs to not only allow security technologies, such as firewalls, for example, to function, but also needs to include direct support for network level security (p. 18). Direct security support is needed to maintain a robust network, to support accounting mechanisms, which need to be able to verify the origin of packets, and to prevent eavesdroppers from determining traffic flows, a requirement identified by several user communities, including the cable television industry (p. 18). In addition, IPng should include direct protections against denial of service attacks and improper use of network layer resources as well as security support for automatic configuration, to ensure that hosts are receiving configuration information from a valid source (p. 18).

To uniquely identify Internet endpoints and nodes, to allow nodes to change their topological location on the network, and to support certain accounting functions, IPng must assign each network layer object in the global Internet a unique name in addition to a unique address (pp. 18-19). This name may or may not have any relation to network topology, location, or routing (p. 18). A node may also have multiple unique names (p. 19).

Consistent with the RFC process and the IPv4 specification, the IPng specification must be "freely available and redistributable" (p. 19) and cannot include any licensing fees for the implementation or selling of IPng software (p. 19). Like IPv4, IPng must be freely available to researchers and developers (p. 20). This unrestrictive access has been an "essential aspect of the development of the Internet" (p. 20).

Recent support for multicasting has been added to IPv4 and this functionality needs to continue to be provided by IPng (p. 20). However, unlike IPv4, which relies directly on the broadcast capabilities of the underlying network media to implement multicasting, IPng needs to implement multicasting at the network layer, such that multicasting is able to function across all network media (p. 20). It should be possible to create and destroy multicast groups and hosts should be able to join and leave multicast groups as needed (p. 21). Traffic should be routed dynamically as changes to groups occur such that manual intervention should not be required (p. 21).

In order to support new user applications such as multimedia and real-time applications and to provide commercial users with a means to receive a particular level of service for a particular cost, IPng needs to allow the network to "associate packets with particular service classes and provide them with the services specified by those classes" (p. 22). This requirement "includes features such as policy-based routing, flows, resource reservation, network service technologies, type-of-service and quality-of-service and so on" (p. 23). Failure to support new applications that require service guarantees may result in users selecting protocols other than IPng

(p. 23). IPng's support for network service should be secure and it should prevent hosts from disrupting the service and resource guarantees granted to other hosts (p. 23). Users should also be able to confirm that they are in fact receiving the level of service that has been guaranteed to them (p. 23).

IPng must also include support for mobile hosts, networks, and internetworks (p. 24). Highly mobile hosts, such as portable devices, as well as less mobile hosts, such as hosts that are unplugged from one network and plugged into another, need to be supported by IPng (p. 24). In addition, mobile internetworks, such as US Navy aircraft carriers, themselves mobile networks, which carry and interact with mobile networks in the form of aircraft must also be supported (p. 24). In order to be able to support civilian airliners, passenger ships, and trains, similar requirements must be met (p. 24).

Like IPv4, IPng should include facilities for testing and debugging networks (p. 25). However, unlike the current implementation of ICMP for IPv4, IPng error handling should include better support for error determination (p. 25). Specifically, error messages should include as a minimum the entire IP header as well as the entire TCP header of the packet that resulted in the error (p. 25).

Lastly, IPng should include support for both IP-based and non-IP-based private internetworking (p. 25). In the case of the former, private IP-based internetworks are used by the research community to develop new services and functionality (p. 25). Private IP-based networks support the extension of the Internet by providing users with a non-disruptive mechanism that may be used to experiment with new functionality. In the case of the latter, current users rely on IPv4 to connect to non-IP islands that utilize other protocols, such as CLNP or AppleTalk, for example (p. 25). These users and uses must be supported by IPng.

Technical Code Summary

- Available as soon as possible
- Simple IP connectivity
- Enable other protocol layers
- Decentralized control and policy
- New capabilities relative to IPv4
- Algorithm and data structure independence
- IPng version interoperability
- Extensible header
- Support for new packet types
- Scalable addressing
- Scalable locally-unique addressing
- Scalable routing
- Hierarchical addresses
- Flexible topology
- Low latency
- High performance subnetwork support
- Processor optimized header
- Robust
- High reliability
- Transition autonomy
- IPv4/IPng interoperability
- Heterogeneous subnetwork support
- ATM support
- Optimal transmission size determination
- Datagram based
- Link sharing with other protocols
- Automatic configuration
- Security support
- Security support, direct
- Accounting support
- Mobile hosts
- Unique names
- Open access
- Free of licensing fees
- Multicasting
- Resource reservation
- Source routing
- Policy routing
- Service selection
- Predictable service for predictable cost
- Mobile networks
- Mobile internetworks
- Multi-homed host support
- Error reporting
- Support for other protocols
- Support for multiprotocol transport networking

Expectations and Specifications

As illustrated by the above description, most of the participant interests of the relevant social groups were acknowledged to some extent and translated into specific criteria for IPng. These criteria were used as a basis for the evaluation of IPng specification proposals. The evaluation, and the eventual selection of an IPng specification, was published as RFC 1752, *The Recommendation for the IP Next Generation Protocol* (Bradner & Mankin, 1995). IPng was renamed IPv6 in accordance with the naming scheme for Internet Protocol specifications and was further defined and clarified in RFCs 1883 - 1887, *Internet Protocol, Version 6 (IPv6) Specification* (Deering & Hinden, 1995), *IP Version 6 Addressing Architecture* (Hinden & Deering, 1995), *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification* (Conta & Deering, 1995), *DNS Extensions to support IP version 6* (Thomson

& Huitema, 1995), and *An Architecture for IPv6 Unicast Address Allocation* (Rekhter & Li, 1995). RFC 2460, which updates RFC 1883, was also examined with respect to the definition of IPv6.

Based on an analysis of these IPv6 documents and RFC 2460, which updates the initial IPv6 specification, I have summarized the relevant social groups' participant interests, understood in terms of technical codes, in tables 1 through 8. Although tables 1 through 8 are organized based on a categorization of technical code descriptions, this categorization is only intended to improve readability and has no relation to the examined documents.

In terms of table layout, the first column of each table is used to list the technical codes identified via the examination of white papers solicited by RFC 1550 and the technical criteria document, RFC 1726. Technical codes that were identified in RFC 1726, but were otherwise not identified in the white papers were assigned to the IETF IPng Requirements relevant social group. Since RFC 1726 translates the participant interests of white paper respondents, BOF attendees, the IPng Area Directorate, and IETF mailing list participants, it was not possible to determine the relevance of these additional technical codes with respect to these various participants. The IETF IPng Requirements relevant social group was thus added to include any and all of these participants for whom the technical code may have some relevance.

The second column in each table lists all of the relevant social groups for whom the technical code has some relevance. The exact scope of this relevance may be deduced from the detailed descriptions for the relevant social groups presented in this chapter. The third column, labeled "TC," corresponds to the technical criteria document, RFC 1726. That is, an "X" in this column indicates that according to RFC 1726 some form of support for the technical code needs to be provided by an IPng proposal. In essence, this column is an indication of whether or not the technical codes required by the relevant social groups were included in the technical criteria for IPng. The last, rightmost column, labeled "IPv6," indicates whether the technical code was included in some way in the IPv6 specification. An "X" in this column indicates that the technical code was explicitly acknowledged by the IPv6 specification. Similarly, a small "x" indicates that although the technical code was not explicitly acknowledged by the IPv6 specification, the specification does not preclude the code's implementation. Finally, a "-" in the rightmost column is used to indicate that the technical code is neither supported nor unsupported and that IPv6 will need to be implemented and deployed before it is possible to determine whether or not the technical code has been supported. This applies especially to technical codes such as "robust" and "low latency," the satisfaction of which requires actual IPv6 deployment and use.

Again, to guide the reader, the identified relevant social groups may be identified using the following letter codes:

- A. The Defense Modeling and Simulation Community
- B. Unified Routing Architecture
- C. Market Viability
- D. IPng Engineering
- E. IPng Transition
- F. IPng Accounting
- G. Electric Power Research Institute
- H. Cellular Digital Packet Data Consortium of Service Providers
- I. IPng Security
- J. Italian National Institute for Nuclear Physics Network Team
- K. U.S. Navy Radio Frequency Tactical Systems Group
- L. Corporations #1
- M. U.S. Navy High Performance Network Group
- N. Asynchronous Transfer Mode
- O. Multiprotocol Interoperability In IPng
- P. The Cable Television Industry
- Q. Corporations #2
- R. IETF Mobile IP Working Group
- S. Multiple Addresses Per Host
- T. Transport and Network Layer Independence
- U. IETF Nimrod Working Group
- Z. IETF IPng Requirements

Table 1: Addressing

Technical Code	Relevant Social Group(s)	TC	IPv6
Hierarchical addresses	B, D, I, J	X	X
Variable hierarchy levels	B		X
Scalable addressing	G, H, J, L, M, P, R, S	X	X
Scalable locally-unique addressing	K, Q	X	X
Address independent of network connection	G, H, M, R		
Source and destination address access	I		x
User information associated with addresses	R		x
Multiple addresses per host	S		X
Transport layer address independence	T		
Unique names	Z	X	X

Table 2: Routing

Technical Code	Relevant Social Group(s)	TC	IPv6
Datagram based	A, B, H, J, K, M, Q	X	X
Low latency	A	X	-
High reliability	A	X	-
Resource reservation	A, H, K, L, M, P, Q	X	X
Hierarchical routing	B, D		X
Non-hierarchical routing	B		X
Routing label	B		x
Source routing	B, P	X	X
Policy routing	E, H, J, K, L, M	X	x
Route aggregation	G		X
Service selection	G, K, L	X	X
Fast route reconfiguration	M		-
Optimal transmission size determination	O	X	X
Predictable service for predictable cost	L	X	-
Scalable routing	Z	X	X

Table 3: Transition

Technical Code	Relevant Social Group(s)	TC	IPv6
Transition in stages	D		X
Transition autonomy	D, J, K, L, Q	X	X
Transition management tools	D, E, M, Q		X
Transition deadline	J		
Transition via routers only	Q		
IPv4/IPng interoperability	C, E, G, H, L, P	X	X
IPv4/IPng dual stack per host	E		X
IPv4 to IPng address mapping	E		X
Available as soon as possible	G, K, P	X	X
New capabilities relative to IPv4	C, J, Q	X	X
Reduced cost relative to IPv4	C		-
IPng version interoperability	Z	X	X

Table 4: Security

Technical Code	Relevant Social Group(s)	TC	IPv6
Security support	A, I, K, L, M, Q	X	X
Security support, direct	F, H, I, K, M, P, Q	X	X
Security improvements relative to IPv4	E, Q		X

Table 5: Other Protocols

Technical Code	Relevant Social Group(s)	TC	IPv6
Link sharing with other protocols	K, L, M, O	X	X
Upper layer OSI application support	G		x
Based on OSI protocols	G, Q		
Support for other protocols	G, L, O, Q	X	X
Support for alternate transport layers	J		X
Support for multiprotocol transport networking	L		x
Enable other protocol layers	Z	X	X

Table 6: Link and Network

Technical Code	Relevant Social Group(s)	TC	IPv6
ATM support	N	X	X
High performance subnetwork support	L, M	X	X
Heterogeneous subnetwork support	E, J, K, L, M, O, P	X	X
Support for low bandwidth networks	H, K, Q, R		X
Support for high bandwidth networks	K, Q		X
Multi-homed host support	M, O		X

Table 7: Packets

Technical Code	Relevant Social Group(s)	TC	IPv6
Compact headers	K, R		X
Extensible header	O, U	X	X
Fixed size headers	R		X
Fixed location header fields	R		X
Simple equality tests for header processing	R		
Support for new packet types	Z	X	X
Processor optimized header	Z	X	X

Table 8: Miscellaneous

Technical Code	Relevant Social Group(s)	TC	IPv6
Multicasting	A, E, K, M, P	X	X
Automatic configuration	D, G, Q, R	X	X
Accounting support	E, L, P	X	x
Accounting support, direct	F, H, I		x
Decentralized control and policy	D	X	X
Network resource identification	G		x
Transport header access	I		X
Transport layer unchanged	J		X
Mobile hosts	K, L, M, P, Q, R	X	X
Mobile networks	K, M	X	x
Mobile internetworks	K, M	X	x
Support for network management	M, P		X
Commercially viable	M		-
Support for time synchronization	M		X
Error reporting	M, O	X	X
Support for international commerce	Q		
Integration with international standards	Q		
Algorithm and data structure independence	U	X	X
Simple IP connectivity	Z	X	X
Flexible topology	Z	X	X
Robust	Z	X	-
Open access	Z	X	X
Free of licensing fees	Z	X	X

Accommodating Programs and Anti-Programs

The solicitation of white papers initiated by the IPng Area via RFC 1550 may be understood as a particular program that strives to include and enrol as many actors as possible. In contrast to a program that aims to exclude certain actors and include others, the program of the IPng Area is a program of inclusiveness. According to Bradner and Mankin (1995), the call for white papers was "intended to reach both inside and outside the traditional IETF constituency to get the broadest possible understanding of the requirements for a data networking protocol with the broadest possible application" (p. 9). This type of program implies that the resulting technology must cater to everyone's needs, that is, so that each group takes up the technology as an unproblematic black box, while simultaneously being general enough so that it is able to enrol as many groups as possible. This is a program of balance between specificity and generality. The saying, "Trying to be all things to all people" aptly summarizes this type of program.

As the rightmost column of tables 1 through 8 indicates, the IPv6 specification is largely based on the technical codes of the identified relevant social groups and appears to balance well between specificity and generality. With the exception of explicit support for mobility and the selection of an OSI protocol as the IPng, IPv6 caters to most of the needs of the identified groups. Specifically, with the above exception, all of the technical codes implied by three or more relevant social groups are either supported in some form by IPv6 or may be implemented given the IPv6 design. These technical codes include:

- Hierarchical addresses
- Scalable addressing
- Datagram based
- Resource reservation
- Policy routing
- Service selection
- Transition autonomy
- Transition management tools
- IPv4/IPng interoperability
- Available as soon as possible
- New capabilities relative to IPv4
- Security support
- Security support, direct
- Link sharing with other protocols
- Support for other protocols
- Heterogeneous subnetwork support
- Support for low bandwidth networks
- Multicasting
- Automatic configuration
- Accounting support
- Accounting support, direct

In terms of explicit support for mobility, IPv6 would need to support an addressing scheme that disassociates addresses from particular network connections. Although IPv6 does not support this disassociation, it does include support for mobility that is similar to the support provided by IPv4. Likewise, although IPv6 is not based on the OSI Connectionless Network Protocol, it nonetheless supports upper layer OSI applications and CLNP. In the case of the former, an OSI compatible transport layer may be used to run above IPv6. In the case of the latter, CNLP may be tunnelled via IPv6. Although IPv6 is not directly based on CLNP, it still supports CLNP and the applications that rely on it. Thus, not only is IPv6 based largely on the technical codes of the identified relevant social groups, in the few cases where certain technical codes are not included, IPv6 still accommodates these codes, albeit indirectly, to some degree.

With respect to programs that require IPng to exhibit features that appear to run counter to the features required by other programs, IPv6 is able to resolve these conflicts through the use of an extensible header and flexible header options. For example, the program of the cable television industry requires that IPng include explicit support for data traffic accounting, which in turn requires that each packet be authenticated to ensure that the sender and receiver are always correctly identified. Given that packet authentication requires both additional bandwidth and processing overhead, this program may be said to run counter to the program of the U.S. Navy, which requires that IPng run efficiently over low bandwidth networks and require little processing overhead.

By including extensible and flexible header options, IPv6 is able to accommodate both programs. That is, the default IPv6 packet header is compact and with the exception of the source and destination addresses includes little additional information. As a result, it is well suited for low bandwidth and low processing applications as required by the U.S. Navy. By supporting extensible and flexible header options, the default IPv6 header may be extended to include support for additional functions. The cable industry may use these options to support authentication and packet accounting. Hosts and routers that understand and support the additional header options may use them to perform accounting and authentication, whereas hosts and routers that do not recognize the options may simply ignore them.

In addition to accommodating and resolving the conflicting programs of the identified relevant social groups, the extensible header may also be conceived as built-in flexibility that can accommodate future programs. That is, IPv6 explicitly supports its own extension so that additional functionality may always be added to the existing technology by existing and future groups. Thus, the extensible header may be understood as the particular technical code that corresponds to the "all things to all people" program of the IETF IPng Area.

Lastly, although the IPv6 specification appears to accommodate most groups, this accommodation is not without ambiguity. For example, although IPv6 includes a flow and traffic class field to support resource reservation and quality of service mechanisms, it is not clear from the design itself exactly how resource reservation and quality of service mechanisms will be implemented. In the case of resource reservation, IPv6 implementations may use the header fields to support exclusive reservations that use all of the available network resources or non-exclusive reservations that share network resources with other traffic. In part, this ambiguity likely corresponds to a lack of operational experience with certain designs and implementations. But more importantly, this ambiguity enables the IPng Area to keep the relevant social groups interested, since the design does not rule out that the implementations required by particular groups will not or cannot be supported.

Conclusion

In addition to the IETF IPng Area, 21 relevant social groups were identified by means of the document analysis outlined in the previous chapter. By soliciting white papers from the Internet community, the IPng Area was able to interest these groups and allow them to contribute to the definition of IPng. The IPng Area's technical criteria for evaluating IPng proposals and its eventual IPng recommendation incorporated many of the groups' technical codes such that the resulting IPv6 specification is well suited to the enrollment of these groups.

By trying to accommodate the widest possible spectrum of participant interests, the IPng Area's mark on the IPv6 specification may be understood in terms of both the direct inclusion of groups' specific technical codes as well as in terms of the inclusion of an extensible and flexible packet header that can accommodate technical codes that were either not directly supported by IPv6 or that may be required in the future. Although the IPng Area's program of inclusion appears to accommodate everyone, this appearance is only possible as long as network development is divorced from the political economy that structures its practice. In the following chapter, I reflect on this bias, consider the irrelevant groups and interests that were not articulated, and relate the specifics of this case study to our understanding of sociotechnical change.

CHAPTER 5: DISCUSSION

Introduction

Using the previous chapter as a basis, I answer the following research questions in this chapter: Which relevant social groups were not involved in the design of IPng and which participant interests were never articulated? What technical codes influenced the design of IPng? And finally, what can we expect of the next generation Internet and future new media as a result of IPng? To answer these questions I begin with a concise summary of the identified relevant social groups and their participant interests and point out that the state and civil society did not partake in the IPng design process. As a result, I suggest that issues of focus traditionally foregrounded by the state and civil society did not inform the IPng design.

With respect to technical codes that significantly influenced the design of IPng, I suggest that IPv4 played and continues to play an important role in the definition of IPv6. Like IPv4, IPv6 is an open standard that is free of licensing fees. The open standard technical code implies that IPv6 will continue to provide application developers with an open platform that may be used to develop new media.

I also identify the unique contributions of IPv6 implementations, network service providers, and Internet users to the fate of IPv6. In the case of the latter, I suggest that Internet users' expectations regarding gratis information and communication services may not be compatible with certain IPv6 technical codes. Despite this possibility, I argue that the history of media is instructive in this regard and should inform any evaluation of the role played by the Internet user.

Finally, I outline the formal bias inherent in the IPv6 design and reflect on broader issues informed by this thesis. I suggest that the IPv6 design appears to successfully accommodate various programs and anti-programs as long as it is removed from the actual context in which these programs are realized. Considered within the political economy of network development, the design accommodates the few with the necessary resources to participate in network development activities. I conclude by extending the insights of this thesis toward our conceptualization of technical codes and our understanding of sociotechnical change in tightly coupled actor-networks.

Irrelevant Social Groups

In response to the IPng Area's white paper solicitation, 21 relevant social groups responded by tracing their respective actor-networks and translating their participant interests into specific technical codes that should be included in the IPng design. Of these 21 groups, seven can be classified broadly as corporate or commercially oriented (Market Viability, IPng Accounting, Electric Power Research Institute, Cellular Digital Packet Data Consortium of Service Providers, Corporations #1, The Cable Television Industry, Corporations #2). Likewise, three of the groups can be classified as military oriented (The Defense Modeling and Simulation Community, U.S. Navy Radio Frequency Tactical Systems Group, U.S. Navy High Performance Network Group). Finally, of the remaining 11 groups, six can be classified as technology oriented, or oriented towards the support of a particular technology (Unified Routing Architecture, Asynchronous Transfer Mode, IETF Mobile IP Working Group, Multiple Addresses Per Host, Transport and Network Layer Independence, IETF Nimrod Working Group). Taken as a whole, 16 of the 21 identified relevant social groups are oriented around either commercial, military, or technology-specific interests.

Of the remaining five relevant social groups, not a single group may be identified with an academic institution – Italian National Institute for Nuclear Physics aside – or local, state, or federal government. Non-governmental organizations are also notably absent. This absence is relevant since it directly corresponds to the absence of an alternate vision for the Internet. The cable television industry certainly provides a direction for the next generation Internet: packets should be accounted and paid for; pay-per-view information and entertainment should be widely available; network resources should be allocated and reserved; customer information should be protected. The U.S. Navy also has a vision: mission-critical applications require data security and integrity; scarce network resources need to be reserved; priority information should receive priority service; moving aircraft carriers should be able to communicate with moving aircraft.

Although the corporate and military next generation Internet has some resemblance to the Internet of the past, it simultaneously and drastically departs from the Internet of the past and, according to the various commercial and military groups, its inherent problems. Although the state and civil society did not respond to the IETF solicitation, one could easily imagine that the next generation Internet desired by the state, or at least by civil society, would be quite different than the corporate and military next generation Internet. One could imagine the desire for public Internet terminals in libraries such that libraries should not have to carry the cost of packet accounting and payment or need to pass on this cost to the library community. Should low-income users have to write one-sentence emails in order to be able to afford packets? Do all users need high resolution video streams that the cable television industry wants to use for pay-per-view entertainment? Or will lower resolution video streams without resource reservation and prioritization suffice? To the extent that the state and civil society have an interest in public access and civic engagement, it seems conceivable that these groups would support a design that supported these goals. Conceivably, the corresponding Internet Protocol would need to be able to run on a variety of older and newer hardware as well as on higher-capacity and lower-capacity media, since most users will not have the newest hardware or access to a high-capacity link. Similarly, the protocol would need to support anonymous communication to enable participation in controversial debates. The ability to conceal one's identity supports open expression and reduces the risk that participants must undertake in order to express unpopular opinions. Lastly, the protocol would need to provide undifferentiated service to an undifferentiated user community to ensure that certain producers and distributors of content, understood in the widest sense to include all forms of Internet communication, were not favored by the system at the expense of others. That is, the barriers to participation should be low and the system should treat users as equally as possible. In sum, these alternate participant interests translate to certain technical codes that are in direct conflict with corporate and military technical codes.

Although IPv6's extensible header options are able to accommodate existing and future programs without granting preference to certain programs at the expense of others, other design features are not as flexible. In particular, the IPv6 header includes fixed and direct support for the labelling of flows and for the specification of traffic classes. In their specification of IPv6, Deering and Hinden (1995) cautiously note that "[t]here is no requirement that all, or even most, packets belong to flows, i.e., carry non-zero flow labels. This observation is placed here to remind protocol designers and implementors not to assume otherwise" (p. 30). Despite this cautious reminder, IPv6 has direct support for the labelling of flows in order to support resource reservation and quality of service mechanisms. How this feature will ultimately be used will depend on many factors, but it is conceivable that service providers would prefer to sell grades of service to grades of users such that quality of service data traffic would constitute a significant portion of traffic on the Internet. The inclusion of the flow and traffic class fields implies that IPv6 provides both differentiated and undifferentiated network service and supports programs that partition Internet traffic and users into classes. In contrast, a design that did not include these fields would not support these programs.

It is worthwhile to consider the role played by the Internet itself in all of this. By this I mean all of the users, organizations, governments, and corporations that connected to each other in exponential numbers in the latter half of the 1990s as well as all of the investments, interests, policies, applications, hardware, and network connections that contributed to this interconnectedness. In 1994, the year just prior to the initial IPv6 specification, the Internet connected a total of 2,217,000 hosts (Internet Systems Consortium, 2005). In 1996, the year just after the initial IPv6 specification, the Internet connected a total of 9,472,000 hosts (Internet Systems Consortium, 2005). And by 1999 over 43,230,000 hosts were connected to the Internet (Internet Systems Consortium, 2005), which by now was the topic of mainstream conversation and was squarely in the public's imagination. The bursting of the dot-com bubble aside, one is left to wonder how similar or different the IPng design process would have been if it had taken place during this time. Would the same groups show up? Or would the group composition and resulting participant interests be drastically different? Given that the Internet has become increasingly interwoven with daily life, it seems likely that if the IPng design process took place today participation would be more diverse than it was in the early 1990s.

Lastly, although neither state nor members of civil society responded to the IPng Area's solicitation for input in 1994, is it not possible that the concerns of these groups were nonetheless articulated via the various IETF IPng meetings, BOFs, and mailing list discussions? Alternately, given the academic (and military) origins of the ARPANET and the IETF, is it not possible that the IETF community raised concerns regarding a commercial- and military-directed next generation Internet? By exclusively analyzing only RFCs, is it not possible that I have too quickly overlooked IETF meeting minutes and mailing list discussions? This, of course, is a possibility. However, the RFCs published with respect to the IPng technical criteria and specification simply do not include any indication that this was the case. The fact that the IPng technical criteria document was made available for comment to the IETF community prior to publication indicates that IETF members must have been satisfied to some extent with the final, published document. Of course, it is always possible that alternate demands were simply ignored by the authors, but this seems to run counter to the IETF process, which tends to try to, as a minimum, acknowledge divergent opinions. As the detailed description in the previous chapter indicates, the actor-network traced by the technical criteria document does not include any of the elements that would suggests that alternate conceptions or programs were articulated.

Interoperability, Open Standards, and the Future of New Media

Six of the 21 identified relevant social groups indicated that IPng should be able to interoperate with IPv4. For most of these groups, interoperability was mandatory, since a significant installed base of IPv4 was already in place and groups wanted to be able to transition over time such that their IPng and IPv4 networks could interoperate. Both the IPng technical criteria document as well as the eventual IPv6 specification acknowledged these participant interests and included support for interoperability. By including this support, IPv6 is able to simultaneously accommodate the needs of groups who are hesitant or ambivalent towards IPv6 and groups that want to modify and move beyond IPv4. By providing some guarantee in the IPv6 specification that IPv6 will be able to interoperate with IPv4, the IPng Area is able to minimize opposition, since groups that are hesitant to upgrade their IPv4 networks have the flexibility to decide when and if they will transition.

In addition, the open standard technical code may also be traced to IPv4 in particular, itself an open standard, and the IETF more broadly, which has historically supported open standards. As an open standard, IPv6 continues IPv4's legacy as an open application platform that may be used to develop new media. That is, unlike the barriers to research and development at the level of the network, which I discuss in more detail below, the history of new media demonstrates that the barriers to participation at the application level are much lower. New media and Internet applications, including web software, instant messaging, and peer-to-peer file sharing, for example, have been developed by individuals who typically lacked the resources of established software companies. One of IPv4's key affordances is that it provides a data transport service for the Internet. Application developers can simply utilize this service as a building block to support more complex functionality and services. To the extent that IPv6 is also an open standard it will continue to provide application developers with an open interface to the Internet and, like IPv4, will support the development of new media.

IPv4, Inertia, and Expectations

During the 10 years that have passed since the initial specification of IPv6 the IETF standards track status of IPv6 has been updated from "Proposed Standard" to "Draft Standard" and initial IPv6 implementations in operating systems and routers have started to appear. IPv6 test networks and deployments have also started to proliferate. Based on these observations, we can say that IPv6 is becoming increasingly "real," since there is a discernible movement from blueprints and specifications to working prototypes that are being made to interoperate with each other. Despite this movement, IPv6 is still far from being a black-box that is taken up unquestioningly by the actors that it is designed to enrol. There are many examples of technologies that have moved from being less "real" to being more "real" and then back to being less "real" again. For example, the Aramis personal transportation system in France oscillated between "fiction" and "reality" multiple times between 1969 and 1987 before blueprints and test results were eventually filed away in filing cabinets, test tracks were disassembled, and prototypes were laid to rest (Latour, 1996). Examples such as this are a reminder that technology's movement is not always unidirectional. Technology is always equipped with a reverse gear that may take the technology back toward its "fictional" origin.

In order to become increasingly more "real," IPv6's supporters will need to control at least three significant actors: IPv6 itself, network service providers, and Internet users. First, IPv6 supporters will need to tame IPv6's implementation. Despite the elegance and additional functionality promised by the IPv6 design, implementations must be able to interoperate and scale to an Internet that is much larger and complex than the Internet of 1994. Certain solutions work very well on a small scale but fail to perform in larger or more demanding network environments. Specifically, the IPv6 design includes support for quality of service provision, multicasting, and authentication and encryption. Although at the time of IPv6's design, various implementations and tests of these individual features were being performed with IPv4, IPv6 requires simultaneous support for all of these features and it is not clear whether and how implementations will work on a large scale. For example, quality of service implementations that make use of IPv6's flow and traffic class fields are still undergoing experimentation and testing and it is still not obvious which method or combination of methods will be used (Potts, 2002).

In addition to being able to successfully implement the IPv6 design for an increasingly expanding and complex Internet, IPv6 supporters will also need the support of network service and telecommunication providers, who will need to use and deploy IPv6 equipment. Although this may seem obvious, what is noteworthy is the investment, understood in terms of capital as well as equipment and experience, undertaken by network service providers to support IPv4. At the time of IPv6's initial specification, several of the identified relevant social groups expressed concern that network service providers would either not move from IPv4 or would move very slowly toward IPv6 given the existing investments in the IPv4 infrastructure. Suffice it to say, these investments and this infrastructure have grown dramatically since 1994 such that the Internet of today is much more closely wedded to IPv4 than it has ever been in the past.

Moreover, IPv4's predicted lifetime has been extended through the introduction and adoption of technologies such as Classless Inter-Domain Routing (CIDR), which enables more efficient use of the IPv4 address space, and Network Address Translation (NAT) devices, which enable users to share IPv4 addresses among multiple hosts. Unlike the studies that supported the need for IPng and predicted that the IPv4 address space would be exhausted sometime between 2005 and 2011 (Bradner & Mankin, 1995, p. 7), a more recent study by a member of the Internet Architecture Board suggests that the IPv4 address space may not be exhausted until as late as

2047 (Huston, 2003, p. 6). Studies such as these and technologies such as CIDR and NAT weaken one of the key actants used to support the IPng effort. The claim that the Internet is running out of addresses has undergone many trials of strength since IPv6's initial specification. This is not surprising given the large scale investments into IPv4 since the mid-1990s. There are heavy costs involved and the actors who have to bear these costs want to ensure that they are necessary.

To the extent that IPv6's fate is closely linked to the fate of IPv4 address space claims, the direction of these claims toward either the immediate or far-off future may be a good indicator of the likelihood that IPv6 will become "real." If in fact the IPv4 address space is far from being exhausted or if some form of consensus is not reached with respect to IPv4's lifetime and IPv6 implementations fail to deliver on the promises of the IPv6 design, then actors who are needed in order to take up IPv6 and turn it into a successful technology will be indifferent or opposed to its adoption. In the meantime, new subnetwork technologies may be developed that render the IPv6 design obsolete. Although IPv6 does have flexibility built right into its packet header, its static header elements, such as the source and destination addresses, their formats, and the flow and traffic class identifiers may be ill suited to the network technologies of 2035.

IPv4 address forecasts and predictions, such as the one provided by Huston (2003), for example, typically trace an actor-network that includes users and user applications that are characterized as dynamic and unstable. That is, nuanced predictions such as Huston's acknowledge that forecasts work with the assumption that an observed rate of address usage will continue into the future at the observed rate. Of course this assumption may or may not hold, since the rate of address consumption may always decline or increase as a result of changes in the actor-network. If new technologies such as mobile IP devices with globally-unique IP addresses become increasingly popular then the address usage rate may significantly increase, leading to a shorter IPv4 lifespan (p. 7). In contrast, if previously allocated but unused addresses are suddenly returned to the IP address authorities and effectively "recycled," then IPv4's lifespan will be prolonged (p. 6). In sum, if the fate of IPv6 is closely linked to the IPv4 address space, then IPv6 is also closely associated with users and user applications that require unique IP addresses.

In a different but related vein, Internet users have grown accustomed to a particular Internet experience that involves to a large measure gratis information and communication services. This is not to say that the Internet is "free" in any crude sense of the term. In addition to the costs borne by the environment as a result of lethal chemicals and materials that pervade computer hardware, users must also be able to afford the necessary computer equipment in order to connect to the Internet. Likewise, "free" email accounts, like "free" television and "free" radio, more often than not involve the sale of an audience commodity to the highest bidding advertisers. However, despite these obvious costs, Internet users experience the Internet in a particular way and expect to have gratis access to information and communication applications.

This expectation has grown over time as a result of Internet users' experiences with now older applications such as email and the web and with newer applications such as instant messaging and peer-to-peer file sharing. Once they are "on the Internet," users expect to be able to gather information and communicate with others at no additional cost. Given this expectation, it seems worthwhile to ask whether or not Internet users would be willing to pay for each packet or for differentiated grades of service? That is, would today's Internet users move from flat-rate, undifferentiated service to either a pay-per-packet or differentiated service? The answer to this question informs our understanding of the importance of IPv6's quality of service and accounting support. Although network and service providers may prefer to charge by the packet or by grade of service, Internet users may have different expectations. If it is the case that these expectations are in fact deeply ingrained, then the IPv6 difficulties with complex features such as quality of service may be unnecessarily delaying IPv6's implementation and adoption.

There is of course an important caveat that accompanies any expectations of user behavior. The history of media has repeatedly demonstrated that the audience is in fact quite flexible and adaptable to new conditions and demands. Perhaps most relevant to this discussion are the transitions already undertaken by radio audiences and Internet users. In the case of the former, the early use of radio to broadcast university and public lectures, live musical performances, and other forms of "high" culture was eventually and overwhelmingly superseded by lowest common denominator programming laden with advertising, despite the expectations of many at the time who predicted that the radio audience would simply have no tolerance for such programming. Similarly, the early Internet was largely free of commercial content such that communication oriented toward advertising or commerce was feverishly discouraged by Internet users. However, over time, commercial content has pervaded email, the web, newsgroups, and instant messaging and is a persistent element of the Internet of today. Given the flexibility characteristic of audiences and users, one cannot expect rigid and inflexible opposition from current Internet users, who may have certain expectations but whose predecessors have also participated in many transitions in the past.

Formal Bias

The inclusion of extensible header options in the IPv6 specification results in a "formal" bias, understood here as "the prejudicial choice of the time, place, and manner of introduction of a system composed of relatively neutral elements" (Feenberg, 2002, p. 81). That is, unlike a more familiar and obvious "substantive" bias where a non-universal or unequal criteria is applied in order to bias a decision or selection, "formal" bias is seemingly neutral and fair and appears this way as long as it is abstracted from the specific context in which it is manifested. The system of law exhibits a formal bias, since although in principle all are equal and deserve due process, the actual practice of the legal system includes many inequalities such as differential access to legal representation. Likewise, although the education system is formally an equal access system, ethnically biased entrance exams and financial barriers determine who is accepted and who is denied access (p. 81).

Understood in terms of formal bias, the IPv6 design may be said to include a relatively neutral element designed to accommodate existing and future programs. From a technical standpoint, extensible headers appear as a wise decision since they enable future growth and change without dictating the direction and shape of this change. However, in practice this future expansion and change will take place within a political economy where significant barriers to participation exist. Historically, research and development into network protocols has been primarily conducted by universities, through computer science departments, by private research and development labs, and by network equipment manufacturers. Developing and testing network protocols is an expensive activity and requires hardware, network test-beds, facilities, and highly educated people.

Given this political economy, it is reasonable to expect that both private research labs and network equipment manufacturers will continue to undertake research that is oriented toward the realization of revenue generation through future networking technology. Although such programs are able to accommodate other programs at times, they are often in conflict with, and implemented at the expense of, other programs. Notably, in a recently published RFC by the Internet Architecture Board, titled *IAB Concerns and Recommendations Regarding Internet Research and Evolution*, Atkinson and Floyd (2004) plead for non-commercial funding for research and development. The authors note that not only is commercial research oriented toward the generation of a direct competitive advantage but that "the funding source can also affect the content of the research, for example, towards or against the development of open standards, or taking varying degrees of care about the effect of the developed protocols on the other traffic on the Internet" (pp. 3-4).

Given IPv6's formal bias, we may identify the university, and even more specifically the computer science department, which possesses the skills and resources necessary for conducting network research, as a key actor in the development of the next generation Internet. The political and economic orientation of network research groups within these departments will influence the direction of the conducted research, which will either support programs oriented toward further Internet commercialization or will support alternate programs oriented toward alternate goals and values.

Technical Codes, Implementation, and Sociotechnical Change

In addition to contributing to our understanding of the dynamics of Internet development, I would suggest that this thesis also informs existing theory and literature that is concerned with technology in particular and sociotechnical change more generally. To begin with, the methodology adopted for this thesis attempts to put into practice the critical theory of technology by defining a method and a corresponding interpretative framework that is informed by this theory. Although the minimalist methodological principles from actor network theory are kept, the methodology aims to avoid normative implications through the explicit inclusion of an evaluative basis that is informed by the critical theory of technology. The success or failure of this methodology corresponds to the quality of the findings and discussion presented in this and the previous chapter and I have no choice but to leave the evaluation to the reader.

With respect to particular conceptual formulations, I would suggest that this thesis informs our understanding of the notion of technical code. Specifically, in addition to the relation between specific participant interests that give shape to particular but seemingly neutral elements of technological regimes, a second, but complementary relation may also be conceived in the "opposite" direction as when existing and seemingly neutral elements are taken up to serve new or additional participant interests.

The case of best-effort datagram delivery illustrates this relation. In the original ARPANET, best-effort datagram delivery was adopted in part to satisfy the participant interests of the military, who required that the network be able to withstand attacks and significant outages at certain points in the network. Unlike a circuit-based approach, which maintains a fixed circuit between two endpoints, a best-effort datagram approach is not committed to any particular path between endpoints and, as a consequence, is able to route around any outages in the network. This best-effort datagram design is an essential feature of today's Internet despite the fact that the Internet is widely used outside of the military context.

As I have described previously, there are certain groups for whom this best-effort datagram service is insufficient and IPv6 strives to accommodate these groups through the inclusion of header fields that may be used to implement quality of service and resource reservation mechanisms. There is some variability in the implementation of these mechanisms such that certain implementations prioritize and reserve network resources much more aggressively than other implementations. Notably, proponents of approaches that reserve resources sparingly and non-exclusively cite best-effort datagram service as an element that needs to be preserved in order to provide undifferentiated service to an undifferentiated Internet community (Eder, Chaskar, & Nag, 2002, pp. 2-3). The best-effort datagram technical code is reinterpreted by these proponents and used to support participant interests that are far removed from the military context that originally motivated this technical code. Stated another way, a seemingly neutral element of a technological regime is interpreted in such a way that the element is re-contextualized and de-neutralized, but this re-contextualization has no necessary relation to the particular circumstances that produced the element in the first place.

In addition, and continuing with the above example, if we distinguish technology construction in terms of the two separate moments of design and implementation, we can see that politics accompanies both moments. That is to say, politics are not limited to only the design phases of technology construction, but permeate the entire development process. In the case of IPv6, the IPng Area translated multiple and varying quality of service technical codes into a single, ambiguous technical code in the IPv6 specification. The particularities of various groups' quality of service technical codes, which included exclusive resource reservation with "hard" real-time constraints in some cases and non-exclusive resource reservation with "soft" real-time constraints in other cases, were translated into a flow and traffic class field in the IPv6 packet header.

For the IPng Area this translation serves two purposes. First, by including support for quality of service in the IPv6 design, the IPng Area is able to keep the relevant social groups interested in the technology. And second, the IPng Area is able to delay implementation decisions until a later time. However, this delay cannot be indefinite and at some point the ambiguity needs to be clarified and resolved. In the case of quality of service, if the resolution of ambiguity involves, for example, end-to-end, exclusive resource reservation then certain groups may resist this resolution, since it has the potential to significantly compromise undifferentiated network service. In contrast, if the resolution involves non-exclusive resource reservation or exclusive reservation on only a per-network basis, then it is possible that conflict may be avoided. In either case, politics accompanies both the design phase, which selects and excludes technical codes, and the implementation phase, which further defines ambiguous codes. This observation implies that groups are able to strategically intervene during both the design and implementation phase of technology construction. Although certain commitments and directions may result following design selection, intervention during the definition of ambiguous technical codes is still possible.

Lastly, and perhaps most importantly, this case study suggests that sociotechnical change is difficult in tightly coupled and rigid actor-networks. Actants with many stable interconnections are difficult to modify, since other actors have performed much work and invested innumerable resources to ensure that these associations are in fact firm and stable. The Internet and computer networks more generally exhibit this tight coupling, which many of the identified relevant social groups foregrounded as they traced their actor-networks. For example, speaking on behalf of corporations, Britton and Tavs (1994) associate the Internet Protocol with corporate organizational structure, network installations, employee skills, Internet users, budgets, service level agreements, teleprocessing lines, firewalls, roaming users, network links, and protocol suites.

Tightly coupled actor-networks are by no means limited to networking technologies. For example, technologies of the workplace, such as the assembly line or computer work pace software, will often also be tightly coupled to other elements. Struggles over the definition and control of the labor process include struggles over the definition and workings of these technologies. However, even in cases where employers may concede that changes are possible to better accommodate the participant interests and needs of employees, the actor-network itself may be utilized to preserve the existing technology and configuration. A company's organizational structure may be optimized for an existing assembly line, which may be housed in a facility whose interior was custom-designed for the machinery, some of which interoperates with the inventory system, which in turn may be linked with the purchasing system, and so on. Although employers may concede that change is possible, an actor-network may be traced to support continuity and deny the realization of this change. It is difficult to argue with a custom-designed building or an inventory system. In cases such as these, technical expertise plays a crucial and strategic role, since the traced associations and seemingly durable actors need to be tested and subjected to trials of strength. How difficult is it to disassociate or re-associate the existing inventory system with an assembly line? How easily can the custom-designed facility accommodate differently designed machinery? The actor-network can only be de-coupled and made less durable by answering questions such as these. The process is necessary to ensure that technology is provided with the flexibility that is needed for the realization of change.

Regrets and Reservations

The primary shortcoming of this thesis is its inability to detail the dynamics of the IPng design process. That is, the process that is presented is quite static and mechanical and fails to account for many of the conflicts and trials of strength and weakness that inevitably constituted this process. This shortcoming may be attributed to a variety of factors, but two factors seem especially relevant. First, the study was limited to an examination of a subset of RFC documents. And second, the style of address present in the examined RFCs was for the most part depoliticized.

In the first case, the decision to only examine RFC documents and exclude IETF meeting minutes and mailing list discussions from the analysis was justified on the grounds that RFC authors would trace the actor-networks that were made durable through past trials of strength, some of which inevitably took place at IETF meetings and on IETF mailing lists. However, the decision to examine exclusively the selected RFC documents did not take into account the possibility that these documents would be largely void of conflicts and trials of strength. That is, the style of address of most of the examined RFCs was matter-of-fact and authors typically traced their actor-networks with little or no regard for other authors or other actor-networks. Given the length restriction (10 pages) for white papers submitted in response to the IPng Area solicitation, perhaps it is not surprising that authors limited their scope and address to detailing their particular actor-networks.

Although in retrospect it would have been beneficial to examine the RFCs more closely prior to limiting the study solely to their analysis, I doubt whether it would have been possible to examine in any detail the many IETF meeting minutes and mailing list discussions. Based on the time and energy required to complete the RFC document analysis, I do not believe that there would have been adequate time to conduct an analysis of these additional texts.

Future Research

As I hope to have demonstrated in this and the previous chapter, the story of IPng is far from complete. Although the initial IPv6 specification was written over 10 years ago, IPv4 continues to be widely used and deployed today. Nonetheless, IPv6 is still IPv4's official successor and it is important to continue to follow its development in order to understand how the next generation Internet will differ from or resemble the Internet of today. More specifically, I would suggest that a productive approach to further research would include a continued examination of the state of the IPv4 address space and IPv6 deployment as well as an ongoing evaluation of the politics of quality of service implementation in particular and the politics of IPv6 implementation more broadly.

With respect to the IPv4 address space, an understanding of address exhaustion claims and counterclaims is important, since IPv6 provides an obvious remedy to this potential exhaustion. We can expect address exhaustion predictions to become more detailed, complex, and nuanced as supporters and opponents of IPv6 subject each others' predictions to repeated trials of strength. Here, opponents may be understood very broadly as those actors who are indifferent towards or against IPv6 deployment and may include, for example, groups who have a large installed base of IPv4 and do not require any of the additional functionality provided by IPv6.

In addition, actual IPv6 deployment may be used to gauge IPv6's success. As long as IPv6 remains in network test-beds and research labs, its fate remains uncertain. Multiple, interoperable, large-scale deployments that are able to operate under typical Internet conditions will be needed if IPv6 is to become a reality. Exactly this sort of trial of strength is required by the IETF if a standard is to move along the standards track from "Draft Standard" to "Standard" and IPv6 is no exception in this regard.

Finally, it is important to follow the IPv6 implementation and the politics that it entails. Clearly, the direction and approach adopted for quality of service provision has far reaching consequences for the Internet and the future of new media. The introduction of differentiated traffic and the prioritization of certain users and certain applications at the expense of other users and other applications have the potential to modify the Internet and its user base in a marked way. More broadly, an understanding of how IPv6's extensible headers are being experimented with is an excellent indicator of the potential directions that the next generation Internet may take. This understanding is a starting point for intervention when the proposed modifications have consequences for the wider Internet community.

Conclusion

The case study conducted for this thesis attempts to put into practice the critical theory of technology. By means of an analysis of texts that document the design of the next generation Internet Protocol, I have answered the following questions: Who was involved in the design of IPng and what did they understand as problems? Who was absent and what issues were never articulated? To what extent did the involved participants influence the eventual IPng design? And lastly, how did existing Internet technology contribute to the definition of IPng?

Based on the conducted research, I suggest that the Internet Engineering Task Force's IPng Area strove to include as many groups and as many interests as possible. Consequently, the IPng specification accommodates the interests of most of the groups that participated in the IPng design process. Equally importantly, the IPng design is extensible and may be modified in the future. However, despite IPng's seemingly universal accommodation of existing and future interests, actual implementations that take advantage of this flexibility will be grounded within the political economy of network research and development.

Many of the groups that participated in the IPng design were commercially oriented. These groups included the cable television industry, mobile telephone service providers, as well as representatives for large corporations. Groups associated with the U.S. military and groups that focused on IPng support for particular technologies also represented a significant portion of the identified groups. The state and civil society were notably absent from the design process. This absence corresponded to an absence of interests with respect to public and universal access and other issues of focus traditionally foregrounded by the state and civil society.

In terms of specific design features, IPv6 supports undifferentiated network service as well as resource reservation and the provision of grades of service. IPv6 also includes direct support for the authentication and encryption of network traffic, which may be used to protect sensitive information and to associate individual data packets with individual users. This latter mechanism may be used for fine-grained packet accounting.

IPv6 is designed to interoperate with its predecessor, IPv4, which is installed and continues to be used throughout much of the Internet. Like IPv4, IPv6 is an open standard that provides Internet applications with a data transport service over a heterogeneous Internet infrastructure. New media developers will be able to continue to design and implement Internet applications that build on and make use of this service and infrastructure.

Although the initial IPv6 design was first defined over 10 years ago, during this time the Internet, and consequently the IPv4 user base, has grown substantially. Even though IPv6 is the official successor to IPv4, it is not clear whether this succession will take place anytime soon. In the meantime, IPv6 implementations are beginning to clarify and define ambiguities in the IPv6 design. Given the importance and potential of the Internet, this process of definition presents a strategic opportunity for the state and civil society to direct Internet development and atone for past absences. However, in order to take advantage of this opportunity an alliance with publicly-funded network research groups is necessary, since these groups possess the resources and skills necessary to participate in conflicts that will inevitably arise as part of this clarification and definition.

REFERENCE LIST

RFC Index Search Engine. (n.d.). Retrieved May 12, 2005, from http://www.rfc-editor.org/rfcsearch.html

Adamson, B. (1994, August). *Tactical Radio Frequency Communication Requirements for IPng (RFC 1677)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1677.txt

Atkinson, R., & Floyd, S. (2004, August). *IAB Concerns and Recommendations Regarding Internet Research and Evolution (RFC 3869)*. Retrieved June 23, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc3869.txt

Bellovin, S. (1994, August). *Security Concerns for IPng (RFC 1675)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1675.txt

Bellovin, S. (1994, August). *On Many Addresses per Host (RFC 1681)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1681.txt

Bound, J. (1994, August). *IPng BSD Host Implementation Analysis (RFC 1682)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1682.txt

Bradner, S., & Mankin, A. (1993, December). *IP: Next Generation (IPng) White Paper Solicitation (RFC 1550)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1550.txt

Bradner, S., & Mankin, A. (1995, January). *The Recommendation for the IP Next Generation Protocol (RFC 1752)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1752.txt

Brazdziunas, C. (1994, August). *IPng Support for ATM Services (RFC 1680)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1680.txt

Britton, E., & Tavs, J. (1994, August). *IPng Requirements of Large Corporate Networks (RFC 1678)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1678.txt

Brownlee, N. (1994, August). *Accounting Requirements for IPng (RFC 1672)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1672.txt

Carlson, R., & Ficarella, D. (1994, October). *Six Virtual Inches to the Left: The Problem with IPng (RFC 1705)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1705.txt

Carpenter, B. (1994, August). *IPng White Paper on Transition and Other Considerations (RFC 1671)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1671.txt

Chapin, L. (1992). *The Internet Standards Process (RFC 1310)*. Retrieved May 10, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1310.txt

Chiappa, N. (1994, December). *IPng Technical Requirements Of the Nimrod Routing and Addressing Architecture*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1753.txt

Clark, R., Ammar, M., & Calvert, K. (1994, August). *Multiprotocol Interoperability In IPng (RFC 1683)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1683.txt

Consortium, I. S. (2005). *ISC Domain Survey: Number of Internet Hosts*. Retrieved June 22, 2005, from http://www.isc.org/index.pl?/ops/ds/host-count-history.php

Conta, A., & Deering, S. (1995, December). *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification (RFC 1885)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1885.txt

Crocker, S. (1995, April 1). *The Address is the Message (RFC 1776)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1776.txt

Curran, J. (1994, August). *Market Viability as a IPng Criteria (RFC 1669)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1669.txt

Deering, S., & Hinden, R. (1995, December). *Internet Protocol, Version 6 (IPv6) Specification (RFC 1883)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1883.txt

Deering, S., & Hinden, R. (1998, December). *Internet Protocol, Version 6 (IPv6) Specification (RFC 2460)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc2460.txt

DeLaCruz, A. (2000, March). *Request for Comments Summary: RFC Numbers 2500-2599 (RFC 2599)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc2599.txt

Dixon, T. (1993, May). *Comparison of Proposals for Next Version of IP (RFC 1454)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1454.txt

Durand, A., & Huitema, C. (2001, December). *The Host-Density Ratio for Address Assignment Efficiency: An update on the H ratio (RFC 3194)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc3194.txt

Eder, M., Chaskar, H., & Nag, S. (2002, September). *Considerations from the Service Management Research Group (SMRG) on Quality of Service (QoS) in the IP Network (RFC 3387)*. Retrieved June 24, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1399.txt

Elliott, J. (1997, January). *Request for Comments Summary: RFC Numbers 1300-1399 (RFC 1399)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1399.txt

Elliott, J. (1997, January). *Request for Comments Summary: RFC Numbers 1400-1499 (RFC 1499)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1499.txt

Elliott, J. (1997, January). *Request for Comments Summary: RFC Numbers 1600-1699 (RFC 1699)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1699.txt

Elliott, J. (1997, January). *Request for Comments Summary: RFC Numbers 1800-1899 (RFC 1899)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1899.txt

Elliott, J. (1997, January). *Request for Comments Summary: RFC Numbers 1900-1999 (RFC 1999)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1999.txt

Elliott, J. (1997, March). *Request for Comments Summary: RFC Numbers 2000-2099 (RFC 2099)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc2099.txt

Estrin, D., Li, T., & Rekhter, Y. (1994, August). *Unified Routing Requirements for IPng (RFC 1668)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1668.txt

Feenberg, A. (1999). *Questioning technology*. New York, New York: Routledge.

Feenberg, A. (2002). *Transforming Technology*. New York, New York: Oxford University Press.

Fleischman, E. (1994, August). *A Large Corporate User's View of IPng (RFC 1687)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1687.txt

Francis, P. (1994, May). *Pip Near-term Architecture (RFC 1621)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1621.txt

Francis, P. (1994, May). *Pip Header Processing (RFC 1622)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1622.txt

Ghiselli, A., Salomoni, D., & Vistoli, C. (1994, August). *INFN Requirements for an IPng (RFC 1676)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1676.txt

Ginoza, S. (2000, May). *Request for Comments Summary: RFC Numbers 2600-2699 (RFC 2699)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc2699.txt

Ginoza, S. (2000, September). *Request for Comments Summary: RFC Numbers 2700-2799 (RFC 2799)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc2799.txt

Ginoza, S. (2001, August). *Request for Comments Summary: RFC Numbers 2900-2999 (RFC 2999)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc2999.txt

Ginoza, S. (2001, May). *Request for Comments Summary: RFC Numbers 2800-2899 (RFC 2899)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc2899.txt

Ginoza, S. (2001, November). *Request for Comments Summary: RFC Numbers 3000-3099 (RFC 3099)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc3099.txt

Ginoza, S. (2003, December). *Request for Comments Summary: RFC Numbers 3200-3299 (RFC 3299)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc3299.txt

Ginoza, S. (2003, December). *Request for Comments Summary: RFC Numbers 3400-3499 (RFC 3499)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc3499.txt

Ginoza, S. (2003, December). *Request for Comments Summary: RFC Numbers 3500-3599 (RFC 3599)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc3599.txt

Ginoza, S. (2003, February). *Request for Comments Summary: RFC Numbers 3100-3199 (RFC 3199).* Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc3199.txt

Green, D., Irey, P., Marlow, D., & O'Donoghue, K. (1994, August). *HPN Working Group Input to the IPng Requirements Solicitation (RFC 1679).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1679.txt

Gross, P. (1994, December). *A Direction for IPng (RFC 1719).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1719.txt

Heagerty, D. (1994, August). *Input to IPng Engineering Considerations (RFC 1670).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1670.txt

Hinden, R. (1994, October). *Simple Internet Protocol Plus White Paper (RFC 1710).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1710.txt

Hinden, R. (1996, June). *New Scheme for Internet Routing and Addressing (ENCAPS) for IPNG (RFC 1955).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1955.txt

Hinden, R., & Deering, S. (1995, December). *IP Version 6 Addressing Architecture (RFC 1884).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1884.txt

Huitema, C. (1994, November). *The H Ratio for Address Assignment Efficiency (RFC 1715).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1715.txt

Huston, G. (2003). *IPv4 Address Space - How Long Have We Got?* Retrieved June 23, 2005, from http://www.arin.net/newsletter/2003_Third_Qtr.pdf

Kennedy, M. (1997, January). *Request for Comments Summary: RFC Numbers 1200-1299 (RFC 1299).* Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1299.txt

Kennedy, M. (1997, January). *Request for Comments Summary: RFC Numbers 1500 - 1599 (RFC 1599).* Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1599.txt

Kennedy, M. (1997, January). *Request for Comments Summary: RFC Numbers 1700-1799 (RFC 1799).* Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1799.txt

Latour, B. (1987). *Science in action : how to follow scientists and engineers through society.* Cambridge, Mass.: Harvard University Press.

Latour, B. (1988). *The pasteurization of France.* Cambridge, Mass.: Harvard University Press.

Latour, B. (1992). Where are the missing masses? The sociology of a few mundane artifacts. In W. E. Bijker & J. Law (Eds.), *Shaping technology/building society: studies in sociotechnical change.* Cambridge, Mass.: MIT Press.

Latour, B. (1993). *We have never been modern.* London: Harvester Wheatsheaf.

Latour, B. (1996). *Aramis, or, The love of technology.* Cambridge, Mass.: Harvard University Press.

Latour, B. (1999). *On actor-network theory: a few clarifications*. Retrieved September 28, 2004, 2004, from http://amsterdam.nettime.org/Lists-Archives/nettime-l-9801/msg00019.html

Law, J. (1987). Technology and heterogeneous engineering: the case of portuguese expansion. In W. E. Bijker, T. P. Hughes & T. J. Pinch (Eds.), *The Social construction of technological systems : new directions in the sociology and history of technology* (1st MIT Press paperback ed.). Cambridge, Mass.: MIT Press.

Lessig, L. (2002). *The future of ideas : the fate of the commons in a connected world*. New York: Vintage Books.

Mason, J. (2002). *Qualitative researching* (2nd ed.). London ; Thousand Oaks, Calif.: SAGE.

McGovern, M., & Ullmann, R. (1994, October). *CATNIP: Common Architecture for the Internet (RFC 1707)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1707.txt

Partridge, C., & Kastenholz, F. (1994, December). *Technical Criteria for Choosing IP The Next Generation (IPng) (RFC 1726)*. Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1726.txt

Patelis, K. (2000). The Political Economy of the Internet. In J. Curran & M. Gurevitch (Eds.), *Mass media and society*. London, New York: Oxford University Press.

Pinch, T. J., & Bijker, W. E. (1987). The social construction of facts and artifacts: Or how the sociology of science and the sociology of technology might benefit each other. In W. E. Bijker, T. P. Hughes & T. J. Pinch (Eds.), *The Social construction of technological systems : new directions in the sociology and history of technology* (1st MIT Press paperback ed.). Cambridge, Mass.: MIT Press.

Postel, J., & Vernon, J. (1982, November). *Requests For Comments Summary Notes: 600-699 (RFC 699)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc699.txt

Postel, J., & Vernon, J. (1982, November). *Requests For Comments Summary Notes: 700-799 (RFC 800)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc800.txt

Postel, J., & Westine, A. (1984, May). *Requests For Comments Summary Notes: 800-899 (RFC 899)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc899.txt

Potts, M. (2002, February). *QoS: Quality of Service for IP Networks*. Retrieved June 23, 2005, from http://www.ngni.org/qos.htm

Ramos, A. (1998, January). *Request for Comments Summary: RFC Numbers 2100-2199 (RFC 2199)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc2199.txt

Ramos, A. (1999, January). *Request for Comments Summary: RFC Numbers 2200-2299 (RFC 2299)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc2299.txt

Ramos, A. (1999, January). *Request for Comments Summary: RFC Numbers 2300-2399 (RFC 2399)*. Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc2399.txt

Ramos, A. (1999, July). *Request for Comments Summary: RFC Numbers 2400-2499 (RFC 2499).* Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc2499.txt

Rekhter, Y., & Li, T. (1995, December). *An Architecture for IPv6 Unicast Address Allocation (RFC 1887).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1887.txt

Reynolds, J. (1991, December). *Request for Comments Summary: RFC Numbers 1000-1099 (RFC 1099).* Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1099.txt

Reynolds, J. (1991, December). *Request for Comments Summary: RFC Numbers 1100-1199 (RFC 1199).* Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1199.txt

Schiller, D. (1999). *Digital capitalism : networking the global market system.* Cambridge, Mass.: MIT Press.

Simpson, W. (1994, August). *IPng Mobility Considerations (RFC 1688).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1688.txt

Skelton, R. (1994, August). *Electric Power Research Institute Comments on IPng (RFC 1673).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1673.txt

Symington, S., Wood, D., & Pullen, M. (1994, August). *Modeling and Simulation Requirements for IPng (RFC 1667).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1667.txt

Taylor, M. (1994, August). *A Cellular Industry View of IPng (RFC 1674).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1674.txt

Thomson, S., & Huitema, C. (1995, December). *DNS Extensions to support IP version 6 (RFC 1886).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1886.txt

Ullmann, R. (1993, June). *TP/IX: The Next Internet (RFC 1475).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1475.txt

Vecchi, M. (1994, August). *IPng Requirements: A Cable Television Industry Viewpoint (RFC 1686).* Retrieved May 15, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc1686.txt

Westine, A., & Postel, J. (1987, April). *Requests For Comments Summary Notes: 900-999 (RFC 999).* Retrieved May 12, 2005, from ftp://ftp.rfc-editor.org/in-notes/rfc999.txt

Winner, L. (1993). Upon opening the black box and finding it empty: Social constructivism and the philosophy of technology. *Science, Technology, & Human Values, 18*(3), 362-378.