

Where is the future heading?

IPv6: obstacles and solutions of the up-and-coming technologies

Monica Haladyna
haladyna@fas.harvard.edu
Spring Semester 2001
CSCIE-132

Table of Contents

1. Introduction
2. The grand scheme of things: IPv4 and the planet Earth as of today
3. IP addresses: bringing new solutions to old problems
4. An immediate answer for the world today: CIDR [RFC 1519]
 - 4.1 NAT or Network Address Translation
5. IPv6: brief history of the new protocol and its structure
 - 5.1 Transitions, deletions, and additions: version 4 and version 6
 - 5.2 What are the extension headers and what functions do they perform?
 - 5.3 IPv6 addressing method
 - 5.4 Representation of IPv6 addresses
 - 5.5 How will IPv6 routing be possible?
 - 5.5.1 RIPv2
 - 5.5.2 OSPF for IPv6
 - 5.5.3 IS-IS: can it be implemented in IPv6?
 - 5.5.4 BGPv4
 - 5.6 Simultaneous existence of IPv4 and IPv6
 - 5.6.1 IPv6 over IPv4 or the "Dual Stack"
 - 5.6.2 IPv6 tunneling and encapsulation
 - 5.7 How will the addresses be renumbered?
6. How do we know that IPv6 works? 6bone testing
7. New complications will arise with new technologies: MEMS
 - 7.1 How about outer space networks?
8. Conclusions: Is there a happy tomorrow?

1. Introduction

“Progress imposes not only new possibilities for the future but new restrictions.”

Norbert Wiener (1894-1964)

The Human Use of Human Beings [32]

This paper is intended to address the issues surrounding the need for a new protocol named Internet Protocol version 6 (IPv6). There are, nonetheless, endless complexities concerning the technologies of the future procure no simple resolution, and at times adding more complications. There is an overview of the Internet Protocol version 4 (IPv4) and the concerns related to the world population growth that will eventually lead to its decline due to the lack of address space, which is required to keep up with the world. There are other issues discussed, such as why IPv6 is needed, its main functions, and why this is the protocol of choice that will replace version 4. Overall, all the transitions that need to take place in order for IPv6 to be up and running, as well as the modifications in the devices that need to be replaced, such as routers. There is also a brief section on routing technologies and adaptations, tunneling and encapsulation schemes, and other devices such as MEMS [9], which can radically change our view of the world and future. These are some of the most significant themes that we must face today in order to try and foresee the future use of IPv6 and how this will affect the world. However, we must question if these are just immediate solutions or if this will alter positively the technology of the uncertain and unpredictable future.

2. The grand scheme of things: IPv4 and the planet Earth as of today

“Ours is a world of nuclear giants and ethical infants. If we continue to develop our technology without wisdom or prudence, our servant may prove to be our executioner.”

Omar Bradley, General, U.S. Army

As of today, the world has approximately 6,149,949,286 people [15] and is growing at an average rate of 6,350,045 births per month. If we consider the period of one year, that would mean that the world population grows at a rate of about 76,200,540 million per year. The U.S. Census Bureau collected these data starting in January 2000 [19]. By taking these figures into account, this leads to the next topic: the need for a new Internet Protocol that can accommodate an extremely fast growing world, as well as being fair in IP address distribution.

With the above estimates in mind, we now take a look at the current number of people and companies that have access to the Internet every day according to Netsizer, which is a private company whose main purpose is that of measuring the average number of hosts using the Internet every day and also determining the average number of users on a monthly basis. Interestingly enough, they measured, for the month of March 2001, around 116,201 million hosts online worldwide [13]. They also have estimated an augmentation of around 4 million hosts per month receiving IP addresses. These averages are made without considering or taking into account all of the supernetting and subnetting [RFC 950] techniques that are regularly used by companies, universities, or ISPs to extend the present need of attaining more IP addresses. The purpose of this

example is to give a comparison and contrast between the number of hosts that the present Internet can support and the world population expansion rate. Currently, IPv4 with its 32-bit addressing scheme can host up to 2 to the power of 32 (4,294,967,296) addresses that could be available. However, this is not the case, since too many addresses due to poor planning and lack of allocation preparation were not distributed in an experienced manner, moreover, nobody ever expected that in a few years the 32-bit address scheme would be at a point of near depletion [2, 9, 10].

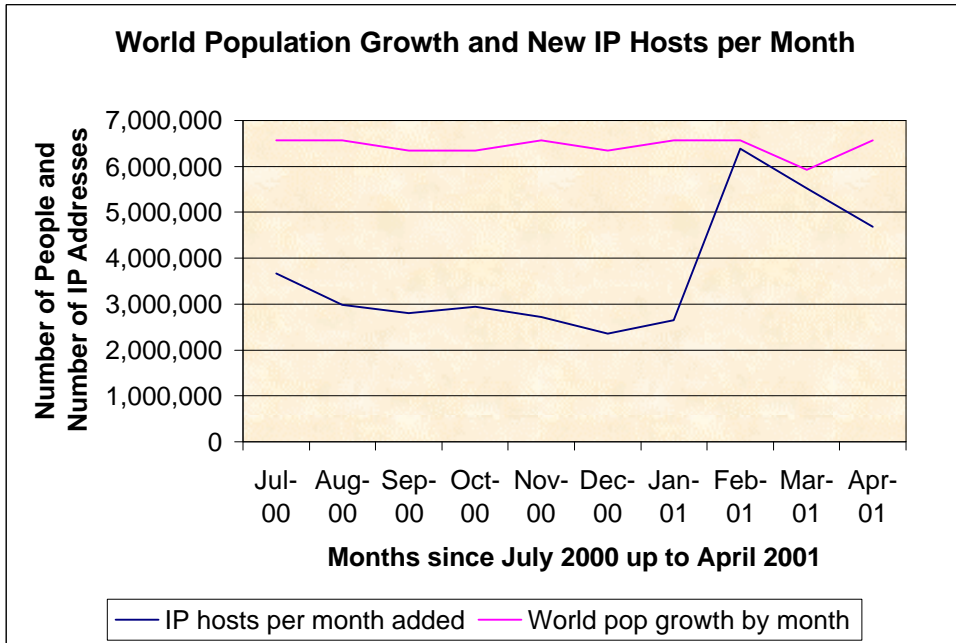


Figure 1. World population growth and IP hosts added to the world on a monthly basis, according to data from the U.S. Census Bureau [19] and the private company Netsizer [13].

In the beginning of the 1980s, when IPv4 was already in use, the IETF came up with a time frame that would project how many networks would need addressing in RFC 1752. They never imagined that by the year 1996, already more than 100,000 networks were running. With IPv4, tentatively, one could have 16.7 million networks and over 4 billion hosts [5], however this did not happen as mentioned previously due to poor address allocation schemes. In 1992 the IETF had already started accepting proposals for a new Internet Protocol that would replace IPv4 in the future; this was called IP the Next Generation or IPng. The IETF had already foreseen the future and decided at this point that action needed to take place by accepting proposals all leading to the establishment of various working groups that would reshape the infrastructure of IPv4 as well as various protocols that will work in junction with the new IPv6 [24].

3. IP addresses: bringing new solutions to old problems

*"He who controls the spice, controls the Universe!"
Baron Harkonnen in Dune (book by Frank Herbert, movie directed by David Lynch)*

Presently, there exist various dilemmas connected with IPv4 that need to be solved expeditiously since the world, population-wise, will not stop growing, and the need for commercial and educational expansion will not cease. The question is: When will we run out of IP addresses? Some calculate that as soon as next year (2002), if we keep doubling the rate of people joining the virtual world every year [2]. One of the complications since the beginning of IPv4 was that classes were assigned liberally without really considering that running out of addresses could eventually develop into a reality. The table below shows a chart of the U.S. top-level domains and the number of hosts that each holds. An important point to consider is that, statistically, the United States has control of well over 75% of all the IPv4 addresses worldwide [4].

com	net	edu	mil	org	gov	us
37245.2	34000.8	78700.30	1837.94	1308.96	940700	18500.30

Table 1. U.S. top-level domains and number of hosts in thousands [18].

Considering the point made above, there is naturally an uneven distribution of IP addresses, and some countries such as India, who just reached one billion inhabitants, have as few as 200,000 IPv4 addresses [14]. Unsurprisingly, many countries are looking forward to a massive transition. At the beginning numerous class B's or /16 as they are known were given out generously to whomever required one. The complication with this is that many of those places fell apart and the addresses were not "recycled." One of the few places in the nation (full class A) that is cooperating in the address "recycling" process is Stanford University [16]. They have proposed that, if there is the need to prolong the survival of IPv4 until IPv6 kicks in, universities and other places where addresses are still abundant but unused, should be obliged to return them with the intention of having them re-allocated to where they are scarce.

4. An immediate answer for the world today: CIDR [RFC 1519]

*"Technical skill is mastery of complexity while creativity is mastery of simplicity."
E. Christopher Zeeman (1925-), Catastrophe Theory, 1977*

There have been other efforts to make use of IPv4 in order to slow down the depletion rate. One of the recent solutions to help aid the need of address depletion as well helping the routing table problem is called CIDR or Classless Inter-Domain Routing [RFC 1519]. What CIDR does is that it is not based on classes (not A, B, or C) but on taking the remaining class C networks of which there are nearly two million left [11, 34] and assigning chunks of these class C addresses depending on how much addresses are needed by an ISP or company. Therefore, if a new ISP needed around 1,000 addresses, they would be given a chunk of 4 consecutive class C addresses, instead of, what was previously done, assigning a complete class B which usually ended up being too much and with many addresses wasted [17, 18]. Thus, those four class C networks would give 1,024 addresses and be sufficient for that particular company.

With CIDR there are some rules applied for the distribution of addresses into four blocks that encompass Europe, North America, Central America, and South America, as well as the Pacific and Asia. It was planned out this way so that each region would have an approximate amount of 32 million addresses. There is, however, one separate chunk containing another 320 million addresses (all class C) reserved for use in the future [17, 18]. CIDR is the result of the effort to procure more time before depletion of all the IPv4 addresses. Nevertheless, CIDR is only a provisional effort to lend a hand to the current world condition, and it is not a future-oriented solution.

4.1 NAT or Network Address Translation

One of the other measures available for alleviating the shortage of IPv4 addresses is by implementing NAT, the Network Address Translation protocol. It works by using some unique addresses, which are set aside for local use. These can be used behind firewalls in large companies to hide the vast amount of users behind the “iron curtain.” It works by having a NAT box where all the people sending requests think that the packets are going just to the NAT box and are unaware that the NAT box will forward all the data to the local user. This in turn decrements the use of having exclusive IP addresses for each individual user in a corporation. Nevertheless, NAT is not an absolute or fixed solution because many of the servers need to have a static IP address and not a dynamic assignment [34].

5. IPv6: brief history of the new protocol and its structure

*“Prediction is very difficult, especially about the future.”
Niels Bohr (1885-1962)*

Since there were so many complexities, the IETF started considering making changes to IPv4 and made the decision to start soliciting proposals on how to make improvements and modifications for a new protocol that could be devised for future implementation. This started happening in the 1990s [5]. The hope was that the new protocol could be formed before IP addresses would run out. One of the first tasks for this new version would be that it could have the ability to support billions of hosts, also making the protocol much simpler by deleting some fields that were never properly used, to somehow find a way to decrement the size of tables used in routing, and enabling the routers to be able to process the packets faster. These are just some of the rudiments that one needed to focus on in order to better the Internet Protocol. Among others, they would want to expand the Quality of Service and make it more applicable for real-time flows, and increase the need for security with the addition of extension headers, which would be optional, and to let the two versions survive together simultaneously.

At this point in time, many working groups began to form, in order for each to take a part in the transition and make this new proposed protocol real. The IETF took action in following this recommendation and giving its approval for this transformation in 1994. One of the groups was called IPngWG or the IPng Working Group. Also the Address Autoconfiguration Working Group and an IPng Transition Working Group helped understand using version 4 addresses in version 6 environments. The IESG was to review and monitor all of the IETF citations with regards to the limitations and implications of IPng, by amending and revising old standards [5, 10, 17, 23].

5.1 Transitions, deletions and additions: version 4 and version 6

“He who loves practice without theory is like the sailor who boards a ship without a rudder and compass and never knows where he may cast.”
Leonardo da Vinci (1452-1519)

Internet Protocol version 6 (IPv6) is the name of the new modified version of Internet Protocol version 4. The quote by Leonardo da Vinci states that practice without theory can be dangerous since it can lead us to random places. Taking this into account, the designers of IPv6 could already correct errors that were made with IPv4, such as having the 8 bits as the Type of Service in the header changed to a diverse type of function called the flow label [5, 23].

32 bits—IPv4 Header

Version (4bits)	IHL (4bits)	Type of Service (8 bits)	Total Length (16 bits)	
Identification (16 bits)		DF/MF	Fragment Offset (13 bits)	
TTL (8bits)	Protocol (8bits)		Header Checksum (16 bits)	
Source Address (32 bits)				
Destination Address (32 bits)				
Options (+)				
Payload (variable length depending on MTU and other factors)				

Table 2. IPv4 header: Highlighted in yellow are the areas that remain the same or with little modification to the field, for example, increased address space. The areas that are in blue have been eliminated from version 6.

32 bits—IPv6 Header

Version (4 bits)	Priority (4 bits)	Flow Label (24 bits)		
Payload Length (16 bits)		Next Header (8 bit ;)	Hop Limit (8 bits)	
Source Address (16 bytes or 128 bits)				
Destination Address (16 bytes or 128 bits)				
Payload (check first MTU, then send data, routers cannot fragment, only the host can)				

Table 3. IPv6 header: Yellow fields remain the same with little or no change. Peach areas are new. Lilac areas have been modified from version 4 and may have a different name and increased function.

Above are the two headers and the changes that took place are highlighted in different colors to make them noticeable. The principal reason for the changes in the IPv6 header is the address bit increase. This was done in order to render a solution to the main problem, that we presently need more addresses in order to accommodate the world population growth and make addresses available to many countries and start-up companies, as well as helping out universities around the world. The transition from a 32-bit scheme to 128 bits is not only necessary to support expansion, but it would give us 2 to the power of 128 IP addresses, roughly 3 times 10 to the power of 38 per every square meter of the entire world [22, 23]. As Huitema [5, 29] calculated, even if a lot of the addresses were completely misused, there would still be around 1,000 IP addresses per square meter, else, if correctly allocated, we could have trillions of them.

The heated debate on deciding which size the IPv6 address should have, came from two propositions: SIPP which supported having 64-bit addresses, whereas the opposition, TUBA, recommended using NSAP addresses that varied in length between 1 to 20 bytes. A final decision was reached that the addresses should be 128 bits long because it could control the growth and not add more layers than the necessary ones [25].

Another transition that took place is that a lot of the options have been removed in IPv6 because, either they have been incorporated into a different field, or were not necessary and just deleted. One such instance is the IHL field from IPv4 that has been used to define the length of the header and, since IPv6 has a diverse approach to this, it was no longer necessary. IPv6 has a fixed header length. One other transition that occurred was to eliminate any of the fragmentation fields: the Don't Fragment, More Fragments and the Fragment Offset; with IPv6 it is already predetermined that all of the routers and hosts need to support packets that are 576 bytes long. In the case that a packet were to be sent and it turned out to be too lengthy, a router along the path would immediately send a message back to the host and tell it that the packet is too long and needs to be fragmented at the point of origin (because routers will no longer fragment). This simplifies things a lot, since fragmentation by routers is quite complicated due to the fact that if a packet were to be lost, the entire packet (fragments and all) would need to be retransmitted. The addition to this is to have an optional extension header allowing the packets to get fragmented [24, 23, 17].

The Time To Live (TTL) field changed to the Hop Limit field. Meaning that one cannot have packets "living forever" and every time that the packet makes a transition from one network to the other the value gets decreased by one (per hop). With version 4 this was the TTL and it was theoretically given in seconds, nonetheless, the routers along the packet's path did not use it this way; instead the routers used it the same as in version 6 by decrementing the number of hops. There is nothing new here, just the name changed.

There is also an increased provision for optional options. This is one of the key differences between version 4 and 6. With version 4 the router(s) would need to take each and every packet, inspect it (all the network level procedures) even if a lot of the options were not even used. For instance most of the time the 8 bits of Type of Service were never used, however, each router would still have the obligation to study each packet when it arrived, and consequently this ended up wasting a lot of time. To not have obligatory options, unless specified, saves a lot of processing time in all the transitions that take place. So with IPv6, if there are no options needed, the router will just skip over this.

Another feature that is essential involves the need for enhanced security. This forms part of one of the options in the extension headers, which are voluntary and not compulsory. If the authentication optional extension header is used, then the sender can be assured that the receiver represents the intended person(s) and vice-versa. The way that this is performed is that the payload gets encrypted using public and private keys using such features as RAS. This is one of the major differences between the two versions.

As for the amplification of Quality of Service, there is a new field referred to as the Flow Label which consists of 24 bits, and has the function of setting up a connection between two end nodes and then decide what will be needed for a specific flow to be transmitted

from end to end. Previous attempts were made with the Type of Service field in IPv4 but, unfortunately, not used very wisely, or for that matter very selectively. Most of the time these bits remained unused, with the exception when they were employed by Diffserv and marking the TOS Byte DS field to give priority to specific types of data streams. While the Flow Label has a more in-depth way of helping devise a less cluttered data network.

There are other fields in IPv6 such as the Priority field which has various sophisticated roles which aid in deciding which packets can congest a network, which can wait, or which need to be transmitted even though some may get lost such as audio and video (in this case no recovery or retransmission). They have a special categorical format that assist the routers in distinguishing which packets are meant for what and if they should have precedence or can wait for a later time.

The true quest of all of this leads to the imperative culmination that IPv6 has many advantages not only structure-wise but has learned from its past history and improved. The core point in all of this discussion is that the header has had a wonderful and evolutionary transition towards simplification and refinement [1, 22, 23, 24].

5.2 What are the extension headers and what functions do they perform?

The extension headers are an optional means of including specific and elective utilities such as needing to have further options like including authentication or any extra data that is disallowed in the main header. These are specified preceding the permanent header and need to follow a certain sequence (note that all of these are optional). The first should be the Hop-by-Hop options header, rendering some information for the routers to process; this is the only type of extension header that the routers take into consideration and actually look within it. The next to follow would be the Routing extension header that could indicate an absolute path from source to destination or a part of the route. Then comes the Fragmentation extension header which dictates how the fragments should be reassembled at end nodes and gives information about how to re-fragment. This is an interesting change since routers would no longer be allowed to fragment, only the source. The next would be the Authentication header that assures the receiver that the particular packet came from the source it claims to have come from. And the last defined extension header is the destination options header which would theoretically be used in case that the end node (receiver) needed to have some type of translation, however, it is not really yet in use [RFC 1883]. It is worthwhile noting that the extension headers are useful not only because they save routers time, since they no longer have to inspect each packet and all the options (as in IPv4), only when the Hop-by-Hop option extension header is present, then it is the only time when it is seen. Extension headers will be a step forward in saving routers processing time because the destination is the one that must process the additional extension headers and with each further it will be solely specified if the next header option is marked. This is the reason why the extension headers need to be in a specific order.

IPv6 Header Next Header = TCP	TCP header and data		
IPv6 Header Next Header = Routing	Routing header Next Header = TCP	TCP Header and Data	
IPv6 Header Next Header = Routing	Routing Header Next Header = Fragment	Fragment Header Next Header = TCP	Fragment of TCP Header and Data

Table 4. Extension headers that are optional make it easier for routers to process [RFC 2460].

5.3 IPv6 addressing method

There are three different types of addresses used with IPv6. The first is Unicast, denoting that it is sent as a detector for a unique interface and, when it is sent to that particular interface, it is identified by that address. A Unicast address is meant for one specific destination [1, 5, 11, 24]. A Unicast address used locally could be depicted as follows:

8 bits	N bits	M bits	P bits
1111 1110	0	Subnet ID	Node ID

Table 5. Single Unicast address in its own network, meaning that there is no prefix yet.

By having a local Unicast address it means that it has only domain in its own network and not in a foreign one. The nodes could have global singularity but their main purpose is to stay within range of their own network; they are not planned for Internet usage, else they would need to request their own prefix.

If the Unicast address were to be used as by a specific provider such as an ISP:

3	N bits	M bits	P bits	125-N-M-P
010	Provider ID	Subscriber ID	Subnet ID	Node ID

Table 6. Unicast address if it were given out by an ISP, considering that the first three bits serve to identify which type of address this is; the provider ID then can allocate the rest to its users.

The next is called an Anycast address and when it is sent off, it is delivered to various nodes who take care of observing which interface is the closest one to deliver it to by taking into consideration where it is heading and the distance it needs to cover. This is done by noticing which destination is the nearest in accordance with the prefix that has tracks of zeros. However, an Anycast global address must not and should not be used as the source address of an IPv6 packet, and it must not be assigned to an IPv6 host but needs to be designated to a router. With IPv6 this very concept of Anycast can be seen as sending to the nearest point/node and, semantically speaking, are very similar to Unicast addresses [5, 24].

The other type is called multicast address, connoting that it has multiple destinations. Multicasting has a special identification permitting a packet to be sent to diverse addresses. This multicast address is used for sending a packet to a chosen group at once and would be denoted with the binary prefix 1111 1111.

8 bits	4 bits	4 bits	112 bits
1111 1111	FLGS	SCOP	Group ID

Table 7. This format would be used for a multicast address.

The binary prefix is 8 bits specified by 1's, then there would be the FLGS, which are set to 000X, and the SCOP measures which values are used for routing. With the SCOP field a 0 it means that it is reserved. 1 indicates that there is an intra-node scope, whereas a 2 depicts an intra-link scope. As for 3, 4, 6, 7, 9, A, C, and D, they remain unassigned for future use [5].

5.4. Representation of IPv6 addresses

IPv6 addresses use fundamentally the same portrayal as version 4 except that the hexadecimal values have been extended to 8 domains of 16 bits each, in contrast with containing 4 parts with 8 bits apiece. The version 6 addresses are represented as in this notation: x:x:x:x:x:x:x having each x represent 16 bits. An example of a Unicast address would look something like this: 198:FA4:893:958:CF:B94:332:954, or others will have many zeros depending on the distribution scheme: 984:0:0:0:4:200:C45:32. In this case, the address having many zeros in each field can be condensed and replaced by a simple colon, which will denote that there are many zeros together. For that matter the address 984:0:0:0:4:200:C45:32 could be written more compact as 984:::4:200:C45:32. When we deal with a multicast address the same rule applies by compressing the sequences of zeros. Also if there is an IPv4 address that is still in use but is being actively incorporated to version 6, then the last 2 hexadecimal fields can be used for the "regular" version 4 address by separating its fields with dots. An instance of this would look like this: 0:0:0:0:FFF:128.5.33.8 and compressed to ::FFF:128.5.33.8 [5, 23]. As of today, the allocation of certain addresses has been already predetermined and others are unassigned for future use as shown in Table 8 below:

Allocation	Prefix (binary)	Fraction of Address Space
Reserved	0000 0000	1/256
Unassigned	0000 0001	1/256
Reserved for NSAP Allocation	0000 001	1/128
Reserved for IPX Allocation	0000 010	1/128
Unassigned	0000 011	1/128
Unassigned	0000 1	1/32
Unassigned	0001	1/16
Unassigned	001	1/8
Provider-Based Unicast Address	010	1/8

Reserved for Geographic-Based Unicast Addresses	100	1/8
Unassigned	101	1/8
Unassigned	110	1/8
Unassigned	1110	1/16
Unassigned	1111 0	1/32
Unassigned	1111 10	1/64
Unassigned	1111 110	1/128
Unassigned	1111 1110 0	1/512
Link Local Use Addresses	1111 1110 10	1/1024
Site Local Use Addresses	1111 1110 11	1/1024
Multicast Addresses	1111 1111	1/256

Table 8. The initial distribution of IP addresses [RFC 1884].

However, there were tremendous debates about the entire version 6 in all the areas imaginable. Should it have a huge allowance in the number of hops or should it be small, or whether or not there should be a 1MB maximum packet size or the 64 KB packet size. The 64 KB packet size was decided on, with the exception that the extension header can permit jumbograms [28]. The most excited dispute was over what size the address should be. They ended up choosing 16 bytes [6,10].

5.5 How will IPv6 routing be possible?

*“The truth of a theory is in your mind, not in your eyes.”
Albert Einstein (1879-1955)*

Currently, a router that needs to send a packet must keep information about all the networks (or connected ones) in its routing tables. These tables can get huge in size and need to constantly maintain sequentially all of the routes that are available for packet transit. What most routers do is keep some default routes that are not included in their own tables, but the problem with this is that the servers (backbones) then must keep all of the possible paths and their tables are enormous. One of the solutions to this is the usage of a tediously designed addressing hierarchy such as with version 6 of IP that considers all of these possible methods and takes into account the routing problems. This implies that there will need to be a wide variety of adjustments in order for the transitions of version 6 to be integrated and slowly emerged worldwide. This process will possibly take a couple of years due to the many infrastructural changes that will need to be tackled and complemented. In certain ways, the routing scheme will be similar to the usage of CIDR, with the alteration of having 128 bit addresses. The protocols that will be used in cooperation with version 6 are mostly called the interior routing protocols. These interior routing protocols (algorithms) have been modified but will be implemented such as OSPF, IS-IS, and a newer version of RIP which is now a version 2. Some of these conversions will be somewhat easier to adjust to because version 6 uses the new extension headers and one of those is explicitly used for routing. With this additional routing extension some very useful functions can be applied such as aiding in host mobility by routing packets to the hosts' present place [5, 6, 7, 23].

Another facet can be to use something similar to “loose source routing” in version 4 and being able to specify which route the packet should take, this way certain nodes along the way can be pre-specified (for various motives such as congestion or skipping certain countries due to security issues). Some of the key factors that were considered for the routing transitions are to be able to have the addressing structure sustain a diverse topological structure and at the same time have as few prerequisites as possible because the address structure and its development at this point is hard to predict. In addition, if the addresses are to be distributed in a hierarchical fashion they should not impose any rules of forcing the routing techniques to be hierarchical as well [25].

5.5.1 RIPv2

For instance, RIPv2 [RFC 1723] is different from RIPv1 because it has the capability of carrying a subnet mask with routing information at the interface level. So if there is the need to have adjustable subnet masks on diverse interfaces in routers they can be configured with RIPv2. Also RIPv2 has the ability to maintain a high rank of security by having verification of password routing information. Additionally, another change that took place from version one is that the next address hop can be decided on for a destination that has already pre-established an advertisement [RFC 1723].

5.5.2 OSPF for IPv6

As for OSPF (Open Shortest Path First) specified in RFC 2740, almost all of the algorithms used in version 4 would remain unchanged except for some elements mostly due to the augmentation of the addresses to 128 bits. The first change would be to have the processing done on a per-link basis and no longer applying to the per-subnet. The elimination of addressing semantics would need to be done, because address labeling will need to evolve into a self-sustaining protocol instead of relying on old specifications and semantics. This would mean that the IPv6 addresses will not form part of the OSPF packets but will be included in LSA payloads and carried by the Link State Update Packets. As well as excluding the network addresses from router and network LSAs and replacing this information with core topology of the LAN (WAN). All of the IDs for routers, link states will need to continue being 32 bits and cannot be assigned as version 6 addresses. The routers that are in the vicinity can be distinguished by Router ID and no longer by an IPv4 address. These are some of the adjustments that are currently being worked on so that when the version 6 is fully operable, the transition will be an obtainable one.

5.5.3 IS-IS: can it be implemented in IPv6?

Intermediate System to Intermediate System Routing Protocol or IS-IS is a protocol that works at the link state level and uses the flooding technique to make link state information available to nodes. At this point there is no current proposal for a version 6, however, due to its very simple nature, it would be useful to continue using it in the future. It has some drawbacks because it was originally designed for the OSI reference model (which has 7 layers) but due to its subtlety it will probably be used in the near future. It uses a “hello” protocol that is flooded to discover its neighbors, as well as subsequent sequence numbers to pass messages. As well as a counter that increases with the messages that are passed along consisting of 32 bits lengthwise, this gives

plenty of time before its inner-counter reaches the limit. As for now it has not yet been set forward for the standards track or a proposal for version 6 been made [11].

5.5.4 BGPv4

Border Gateway Protocol version 4 (BGPv4), at this point in time is considered the principal external gateway protocol, or the ERP (Exterior Routing Protocol), meaning that it is used to route packets that are headed for another autonomous system. BGP has special characteristics unlike other protocols by using TCP to depend on for the transportation of messages to and fro the autonomous systems and the nodes associated with them. This has some superior quality because during the use of TCP, there is a connection-oriented nature that makes it reliable, and there is no need to get new identification numbers or all the other factors that are implied. Once the handshake is done, the rest is taken care of because TCP will use all of its essential communication standards and apply them to BGP. The drawbacks are that BGP is completely adapted to the usage of 32-bit addresses and it will not be able to make a steady transition to settle in with version 6. Therefore another protocol was proposed by the IETF as an external routing protocol and it's called Inter-Domain Routing Protocol or IDRP. The principal reasons for choosing IDRP over BGP is that TCP should not be used explicitly, instead it would be a better idea to use UDP to exchange messages. IDRP is a transition for the future, by accommodating various diverse protocols and having the advantage of utilizing prefixes of flexible lengths. Furthermore, BGP is devised as to maintain a complete record of all the AS which it must pass through, whereas with IDRP a new ideology of aggregating the information is made possible.

5.6. Simultaneous existence of IPv4 and IPv6

One of the foremost vital themes about the transition from IPv4 to IPv6 is how both will need to coexist for a number of years together and manage to be present concurrently. Presently, there are extensive choices of plans available to make this transition a reality before the world runs out of version 4 addresses. For this to be enabled there are certain aspects to be taken into account, such as the need for protocol translators, implementing network address translators, how mobile IP can still be functional without breaking the connection, as well as application layer gateways (ALGs) to be executed.

5.6.1 IPv6 over IPv4 or the "Dual Stack"

Right now IPv4 nodes do not understand IPv6 [RFC 1933] or vice versa making the transition effort of mounting complexity. One of the options is to have the use of a "Dual stack IP" layer, which can provide functionality for both IPv4 and IPv6 at the router level as well as for the hosts. This type of stack is usually referred to as IPv6/IPv4 node and it would be the format in which one could receive and send both IPv6 and IPv4 packets. If interoperation were to take place with an IPv4 node then the IPv6/IPv4 node could use the IPv4 packets. If it were backwards, the IPv6 node could then speak to the IPv6 networking. They would in turn need to have both capabilities in order for them to talk with each other. To learn to distinguish which node supports what type of IP scheme, DNS could be used to give back the IPv6 address in order to determine which nodes are IPv6/IPv4-able and which are not. However, the DNS will only be able to return an IPv4 address. With this "Dual Stack IP" method there are five ways to distinguish what kind of service they will require for each of the existing types of nodes: if it is an IPv4-only node,

meaning that it has no connection with IPv6, or if it is an IPv6/IPv4 node which would understand both IPv4 and IPv6. If the node is an IPv6-only node, which does not understand IPv4, else if the node is both IPv6/IPv4 and IPv4-only are both IPv4 nodes. Or if IPv6/IPv4 and IPv6-only are both IPv6 nodes [7,8].

Type of Node	IPv4 Only	IPv6 Enabled	IPv6 Only	IPv6 Enabled
IPv4-only Node Does not have IPv6 Dual Stack	IPv4-only by using IPv4	IPv4-only by using IPv4	Communication between client and server not possible	IPv4-only by using IPv4
IPv6 client that does not have dual-stack. IPv4 capable. But has one IPv6 node that is configured and has a dual stack that has a connection with one IPv6	IPv4-only by using IPv4	IPv4-only by using IPv4	Communication between client and server not possible	IPv4-only by using IPv4
IPv6 only node which knows client is IPv6 aware. Would need a double stack to communicate	Communication between client and server not possible	Communication between client and server not possible	IPv6	IPv6
IPv6 aware client and IPv6 enabled node. Has only one stack and one address. Fully IPv6.	IPv4-tunneling or proxies	IPv4-tunneling or proxies	IPv6	IPv6

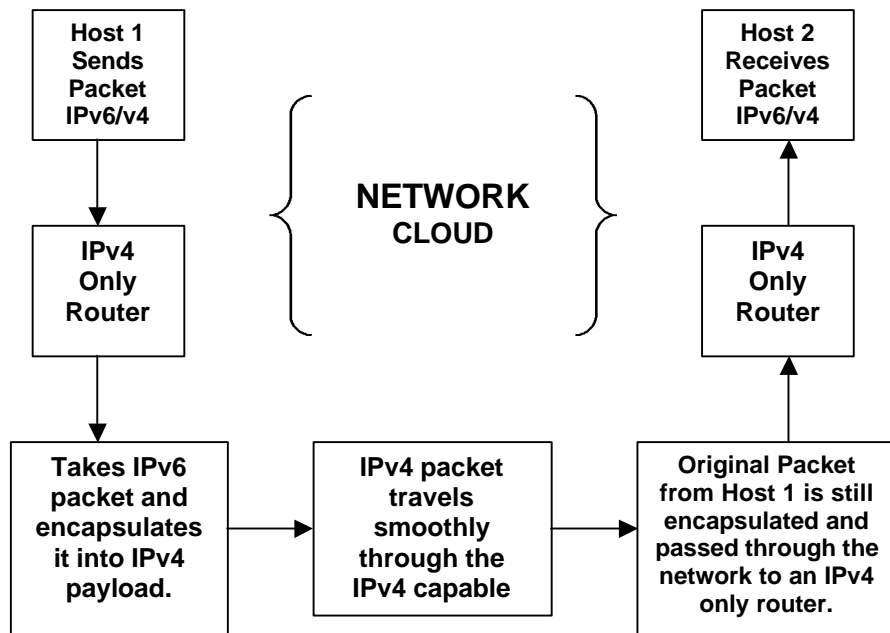
Table 9. Communication between client and servers with IPv4 and IPv6 both enabled, IPv4 or IPv6 only nodes, or IPv6 enabled.

However these are not simple issues to embark upon for the reason that all sorts of transformations would have to be performed. One of the best ways would be to have the ISPs be dual-stack-able by running version 4 and 6 simultaneously, and using version 6 by encapsulating version 4 in it via tunneling techniques. Therefore the routers could still use IPv4 addresses to facilitate the 6 over 4 techniques. Or they could use IPv6 routers concurrently with IPv4 routers and using other processes for tunneling. In order to discover their neighbors they could use the Neighbor Discovery protocol that would let IPv6 hosts have a Domain Name IPv4 prefix address in order to be able to route. Then if there was the case that a version 6 location wanted to send something to a 6 to 4 place in order to get a packet through to the version 6 network they could use the DNS to return an IPv6 and IPv4 address for that host. Then the version 6 host will have an addition to its prefix which will be the 6 to 4 address and will be able to send packets across to its final destination hoping that the receiving end node employs 6 to 4 also. There is another feature named Top Level Aggregation or TLA that specifies that the packet is not local and it is usually a version 6 packet that has been encapsulated into a v4 packet in order to be able to be put through the network(s) [6, 7, 23].

5.6.2 IPv6 tunneling and encapsulation

Another option is to have the IPv6 run over IPv4 through tunneling by encapsulating IPv6 packets inside the IPv4 headers (such as in mobile services which use IP-IP) and

putting the packets through regular IPv4 networks. With tunneling there would be the option of having it be either configured or automatic [7]. As put in RFC 1933, the most important thing is to have interoperability during the transition period for the base of IPv4 functionality. With this technique, the IPv6 sender's node would encapsulate the packet into an IPv4 payload and could then have the packet addressed to the IPv6 receiving node at the end of the tunnel. If along the way the packet passes through many IPv4-aware routers, these will not know that the entire IPv6 packet is encapsulated until it arrives at a IPv6-aware environment where it will then be decapsulated and sent to the final IPv6 node. With configured tunneling the IPv6/IPv4 happens when the IPv4 final destination address has been determined by pre-arranged information on the encapsulating node [6, 23]. As with Automatic tunneling, the endpoint address of the tunnel has been already determined by the IPv4 address encapsulated in the destination address of the IPv6 packet. Some of the related concerns related to tunneling would be with regards to the MTU or Maximum Transmission Unit size because if it were the case where IPv6 was being used, some of the links along the way cannot handle packets that exceed the specific size with IPv4 links. Thus, fragmentation could take place since the fragmentation fields in the IPv4 header are still present even if they had been set to the Do Not Fragment bit. This is also one of the reasons why manual tunneling would come in handy. These are just a couple of examples of issues dealing with the complexity of 6 over 4 presently.



5.7 How will the addresses be renumbered?

The transition could take place by first adding the new addresses to the routers and letting the old addresses fade gracefully. Also announcing the new prefix information in the router advertisement messages and at the same time to broadcast on a short term basis the old prefix with a limited amount of time or to altogether stop announcing the old

prefix. To handle routing one would try to advertise the new routing information to other nodes and not stop advertising the old routing information. There would need to be as well a new updated version of DNS tables in order for reverse lookups to take place and be able to handle both new and old addresses. Then after some time when most of the systems have been upgraded to IPv6, have the DNS tables dissipate with the old version 4 address information and stop advertising the routing information only the new data. Then they could change all the old addresses from the routers, cancel the version 4 reverse lookups and perhaps implement something like an automatic router renumbering scheme in order to smoothly make the transition from one to the next [1, 30].

6. How do we know that IPv6 works? 6bone testing

In order for IPv6 to be of any value besides the theoretical corrections and modifications that took place, there needed to be some way of testing, experimenting, and checking that this protocol is in fact one of the answers that will lead the world to a more sustainable form of stability [6]. The 6bone has been tested in various “test beds” in order to detect any flaws or mistakes that could turn out to be disastrous if they were to be deployed worldwide without a thorough examination. 6bone is as of now a worldwide project that is being supervised by NG-trans, which is an IPv6 related transition group from the IETF [20]. There are currently, as of May 7, 2001, a total of 811 6Bone sites worldwide [20]. 6bone simply stands for IPv6 backbone and in the beginning 6bone started by using IPv6 over IPv4 in tunneling and encapsulation techniques and testing it over IPv4 networks to see what the results would be like. One of the innovative 6bone sites was in Japan; it is called the WIDE 6bone project which started in 1995. This is presently one of the biggest sites for IPv6 testing. This innovative project which was IETF approved [1], pointed out that the rigorous testing led to the conclusion that address selection (hierarchically) and renumbering methods are just about the most important factors if the dual stack is running in a selected environment. It led to the conclusion that the abolishment of IPv4 would need to take place in a “gentle” manner, by slowly merging IPv6. Once the routers, nodes, and DNS got their tables updated, then IPv4 can begin to fade away slowly to give way to version 6. Other speculations which they noticed in the WIDE 6bone was that there needs to be a way to guarantee that packets will be delivered even though some may need to be transferred via an IPv4 link. As well as having their outgoing connections from a location be pre-set by an automatic translator which can be set up at the 6bone (backbone). There are however many questions and solutions which still need to be worked on. Around the world there is a vast total of 811 6Bone sites, of which many are connected using a world 6bone. Most of the regional 6bones are capacitated by running IPv6 over IPv4 via tunneling techniques. The 6bone operates under the IPv6 Testing Address Allocation, which is specified in RFC 2471.

In order for an individual to get connected to a 6bone there are some standard rules to be understood. First, since the Globally Addressable IPv6 has a three level hierarchical structure which first includes a Public Topology, then a Site (location) Topology, and an interface identifier, which can be a 64-bit number that is singular in a LAN. The Public Topology has two aggregators, usually called the Next Level aggregator, which can be an ISP, and the second is named the Next Level Aggregator that can be the end location. Then the ISP gives the end location (user) their prefix, which gives them IPv6 access to their backbone. However, the Top Level Aggregators are assigned by the 6bone collaborators; these would be the first 16 bits. The RIR or International Regional

Internet Registry would give these out. To get an IPv6 address first find an ISP that uses IPv6 and then one would get a prefix from them. Or another way would be to use 6 over 4 automatic tunneling. There are some apparatus that will need to be replaced such as routers and having as well a device that can be a host. There are some companies that already have fabricated and tested such equipment such as Bay Networks, 3Com, Hitachi, Nokia, Telebit, Sumimoto and Digital [20]. Another suggestion, once that the router has been configured to support IPv6, is to be able to forward packets through the diverse interfaces, hopefully using Ethernet [27], to the end-user stations. Next would be to find a 6bone to connect to with the locations' TLA. The administrator could build an IPv4 tunnel from the site using an IPv6 router and in turn having this one joined to the 6bone in order to dispatch the packets. There is a lot of information about the TLA 6bone test employment in RFC 2471.




Table 10. The 6bone sites around the world. Image from reference [34].

Once this is completed, there is plenty of software that can be downloaded for free [33] and that can be used for updating networks and their components. Above is an image that displays many of the 6bone networks that are up and running throughout the world [34, 3, 6]. For a full list of countries that are building up 6bone sites look at [20] as well as a chart that is designed on a daily basis according to who is joining 6bone [21, 31].

7. New complications will arise with new technologies: MEMS

"In the search for truth there are certain questions that are not important. Of what material is the Universe constructed? Is the Universe eternal? Are there limits or not to the Universe? ... If a man were to postpone his search and practice for Enlightenment until such questions were solved, he would die before he found the path."

Buddha

MEMS are Microelectromechanical Systems (MEMS) containing miniscule mechanical elements that are connected with electrical processors by circuits [26]. MEMS are measured in microns, meaning millionths of a meter and are so tiny that they can be compared to the width of a human hair. These devices are fabricated similarly to microchips, and have the properties of being able to detect heat, movement and other factors. Currently DARPA has a lot of research going in this area, from having miniscule robots to perform surveillance, to having them injected in a person's blood stream to detect diseases. If MEMS really take off (right now they are very expensive for mass production) and were to be given such properties as IPv6 addresses, then the question is, will there be enough IPv6 addresses left in 20 years? These devices would have such a wide range of capabilities, and considering that they are practically invisible (some are, some are not), and with certain "intelligence" to perform certain tasks [9, 12], do we run into the same problem of IPv4, will IPv6 suffice? IPv6 offers 3 times 10 to the power of 38 singular addresses meaning that if each of these MEMS were to be identified, it is comparable to having every star in the Milky Way labeled as well as Andromeda (the next galaxy) and all of its stars. However, with all the technological craze, perhaps, it will not be enough.

7.1 How about outer space networks?

Glimpsing at the future permits us to think not only about Earth related boundaries, but this all leads to the question of outer space. In the near future with the evolution of all these technologies, it would not surprise me if the outer space networks would utilize IPv6 addresses for every and each component on their stations. This would be quite practical if they had addresses for every space suit, computer, radio, etc. and could monitor each item from the Earth. If here on Earth they would plan to have microwave ovens with IPv6 addresses why not have IPv6 addresses for more important components in outer space? These devices could be monitored all the way from Earth to check if there are any problems and one would perhaps even be able to fix them from here. The NASA websites don't mention IPv6 in any of their programs, so this is just an idea that could perhaps be implemented in the future.

9. Conclusions: Is there a happy tomorrow?

"Technological progress is like an axe in the hands of a pathological criminal."

Albert Einstein (1879-1955)

There is really no way to predict how the future protocol IPv6 will work. There are so many things to consider with the transitions that need to take place presently, which is already overwhelming to even think about, being worse for first world countries such as the United States who must undergo tremendous changes in their entire infrastructure. Other countries will have it a little bit easier since some are not as developed and will begin buying the routers and equipment that is compatible with the newer version.

Ethically, there has to be some sort of mediation and international agreement on how all these new technological advances will affect humanity. There needs to be logic, thought, good intentions, and morality when dealing with high-tech issues. This is because technology needs to have regulation, especially if we are going to be expanding to astronomical levels of IPv6 addresses, however, this is not only in IP terms but in human population growth, problems are foreshadowed. IPv4 has mainly been used for the first world countries, whereas IPv6 will be accessible to anyone who can afford to buy a computer and have access to a telephone line, satellite connection, fiber optics, DSL, or Cable. This will render third world countries the opportunity to integrate themselves into the virtual world since many of these countries have been complaining that not sufficient addresses are available for them, and living in a capitalistic, democratic society this is the only way to bring fairness into the world. This brings us to another point: will this change to IPv6 suffice? Perhaps if MEMS take a "life" of their own, IPv6 addresses will exhaust at unseemly speed, as well as assigning IPv6 addresses to every sector and device aboard a space mission. With MEMS [12] there could literally be billions of them to carry out tasks that humans can handle as well. Another worry is who will patrol the world? Perhaps there will be now more hackers and crackers than ever before despite the authentication and encryption offered by version 6.

There are many topics which had to be excluded in the paper that were very interesting, but this is fundamentally an overview of the changes that are taking place on a worldwide scale because IPv4 addresses will at some point (perhaps 10 years from now) become obsolete with the New Generation IPv6.

Acknowledgements

Thanks to Professor Scott Bradner (sob@harvard.edu) for his class and the opportunity to develop this paper and for his rigorous and comprehensive notes that have served well as guidelines in exploring networking topics.

To the teaching assistants, Mark Gaynor (gaynor@harvard.edu) and Conrad Nobili (nobili@harvard.edu), who have always been extremely encouraging, prompt, helpful, and excellent guides.

References

[1] Yamamoto, Kato, and Sumikawa, "Deployment and experiences of WIDE 6bone," <http://www.v6.wide.ad.jp/Papers/yamamoto/>

[2] S. Mace, "New protocol not a root canal or a forklift upgrade," <http://www.Stardust.com>

[3] Rockell and Fink, "6Bone backbone routing guidelines," RFC 2772, <ftp://ftp.isi.edu/in-notes/rfc2772.txt>

[4] C. Krishnaswamy, "IPv4," <http://www.timesofindia.com/090101/09mban58.htm>

[5] R. Jordan, "IPv6: the future of the Internet Protocols," http://www.eece.unm.edu/faculty/rjordan/595-025/jeff/ipng_paper.html

- [6] Fische and Heissenhuber, "Mobile IPv6: mobility and support for the next generation," IPv6 Forum,
<http://www.IPv6forum.com/>
- [7] Slattery and Terry, "A look at the future: IPv6," The Network Monitor,
http://www.ccci.com/product/network_mon/tnm34/IPv6.html
- [8] Carpenter and Moore, "Connecting IPv6 routing domains over the IPv4 Internet," IBM and University of Tennessee,
http://www.ieng.com/warp/public/759/ipj_3-1/ipj_3-1_routing.html
- [9] Information on MEMS at DARPA, wide range of projects,
<http://www.darpa.mil/>
- [10] J. Raj, "Chapter 18: IP next generation (IPv6),"
<http://www.cis.ohio-state.edu/~jain/>
- [11] Myelabs.com staff, "We are running out of IPs!!!,"
<http://www.myelabs.com/reviews.asp?name=1242>
- [12] MEMS,
<http://www.lcs.ece.cmu.edu/research/MEMS/>
- [13] Netsizer, "Internet growth reports,"
<http://www.netsizer.com/>
- [14] C. Krishnaswamy, "IP version 6 to take off on a transcontinental mode,"
<http://www.timesofindia.com/090101/09mban58.htm>
- [15] D. Levine, world population Java applet,
<http://www.ibiblio.org/lunabin/worldpop>
- [16] M. Duffy, "Stanford move rekindles Net address debate,"
<http://www.nwfusion.com/news/2000/0124IPv4.html?nf>
- [17] J. Lawrence, "The future is ahead of schedule," Caltech/MIT Enterprise Forum,
<http://www.trillium.com/>, http://www.its.caltech.edu/~entfor/f_video/1119165.17.pdf
- [18] Netsizer, "Current stats for top level domains,"
<http://www.netsizer.com/daily/TopLevelDomain.html>
- [19] U.S. Census Bureau, "World POPClock Projection,"
<http://www.census.gov/cgi-bin/ipc/popclockw/>
- [20] 6Bone registry by countries,
http://www.6bone.net/6bone_countries.html
- [21] 6Bone chart designed daily from statistical growth and information,
<ftp://ftp.isi.edu/6bone/6bone.db.gz>

- [22] More information on IPng Working Group from Sun Microsystems,
<http://playground.sun.com/pub/ipng/html/ipng-main.html>
- [23] S. Bradner, "IP next generation: a path to the future,"
<http://www.sobco.com/e.132/reading/IPv6.html>
- [24] Deering and Hinden, "Internet Protocol, version 6 (IPv6)," RFC 2460,
<ftp://ftp.isi.edu/in-notes/rfc2460.txt>
- [25] Wessman and Petri, "Existing routing protocols and IPv6,"
<http://www.tml.hut.fi/Opinnot/Tik-110.551/1996/ip6routing.html>
- [26] Hui and Elliot, "Microelectromechanical Systems,"
<http://bsac.eecs.berkeley.edu/~elliot/mems.html>
- [27] Crawford, "Transmission of IPv6 packets over Ethernet networks," RFC 2464,
<ftp://ftp.isi.edu/in-notes/rfc2464.txt>
- [28] Borman, Deering and Hinden, "IPv6 jumbograms," RFC 2675,
<ftp://ftp.isi.edu/in-notes/rfc2675.txt>
- [29] Crawford and Huitema, "DNS extensions to support IPv6 address aggregation and renumbering," RFC 2874,
<ftp://ftp.isi.edu/in-notes/rfc2874.txt>
- [30] Crawford, "Router renumbering for IPv6," RFC 2894,
<ftp://ftp.isi.edu/in-notes/rfc2894.txt>
- [31] Software implementations,
<http://playground.sun.com/ipng/ipng-implementations.2.html>
- [32] Quotes by Norbert Wiener,
http://www-groups.dcs.st-andrews.ac.uk/~history/Mathematicians/Wiener_Norbert.html
- [33] 6bone GIF drawing of all the countries and networks connected worldwide,
<http://www.cs-ipv6.lancs.ac.uk/ftp-archive/6Bone/Maps/backbone.gif>
- [34] Sun Microsystems, "IPv6 and the future of the Internet,"
<http://www.sun.com/software/white-papers/wp-IPv6/>