

Evaluating IPv6 in Peer-to-Peer Gigabit Ethernet for UDP using Modern Operating Systems

Burjiz Soorty

School of Computing and Mathematical Sciences
Auckland University of Technology
Auckland, New Zealand
burjizsoorty@hotmail.com

Nurul I Sarkar

School of Computing and Mathematical Sciences
Auckland University of Technology
Auckland, New Zealand
Nurul.sarkar@aut.ac.nz

Abstract— Most of the previous works on evaluating IPv6 have focused on network throughput and packet delay performance studies. A thorough literature review reveals that a very limited research have been conducted on evaluating IPv6 using quality of service (QoS) parameters and modern operating systems (OSs). The objective of this paper is to evaluate the UDP performance of IPv6 in a peer-to-peer Gigabit Ethernet network using the latest four different MS Windows and Linux-based client and server OSs. The QoS metrics such as throughput, end-to-end delay, jitter, and CPU utilization are measured empirically to determine which OSs provide the best bandwidth performance over IPv6 networks. The impact of packet length on system performance is also reported. Results show that for IPv6 Linux network using Ubuntu 10.04 performs significantly better than Windows 7.

Keywords: Bandwidth, IPv6, operating systems, packet length, UDP

I. INTRODUCTION AND BACKGROUND

Network performance is one of the most critical issues that need to be taken into account when planning, configuring, optimizing and upgrading networks as well as during the design, deployment and development of network infrastructure. It is therefore imperative to measure network performance, gather detailed analysis of its end-result, and thereof produce new research to help network professionals and engineers to achieve better performance on current technology and standards. A study on measuring key performance metrics such as network throughput, packet delay (time for a packet to send from a source to a destination), packet jitter (i.e. delay variance) and CPU utilization (i.e. CPU processing resources consumed by the host to transfer data across the network) need to be critically assessed in the evaluation. Furthermore earlier research has shown that factors such as operating systems, type of protocol and packet length affect these parameters [1]. Earlier studies have evaluated the performance of IPv4 and IPv6 for TCP. However a very limited research has been done using modern operating systems for UDP (transport-layer protocol for delay-sensitive traffic).

Rapid advancement in technology, increasing IP-based appliances, and growing use of IP-based applications such as delay-sensitive VoIP (Voice over Internet Protocol), IPTV (Internet Protocol Television), Mobile Internet, Video-

Conferencing and Online-Gaming all have invariably led to the exhaustion of the IPv4 address pool [2].

As this shortage of IPv4 addresses nears exhaustion and deployment for IPv6 technology grows, organizations are left with four major choices. One option to continue using IPv4 in their internal networks and use NAT (Network Address Translation) to communicate across the current IPv4 infrastructure and use NAT-PT (NAT protocol translation) upon the soon-to-be predominant IPv6 internet [3]. This option although foregoes any upgrade costs for an organization to upgrade to IPv6 it does also hinder progress towards establishing a predominant IPv6 internet.

The second option is to perform a complete internal-network migration to IPv6 and use tunneling mechanisms such as 6to4 GRE (Generic Routing Encapsulation) [4] and ISATAP (Intra-Site Automatic Tunnel Addressing Protocol) [5] which allow IPv6 communication across current IPv4 infrastructure.

The third option also allows an organization to continue using IPv4 and at the same time gradually upgrade to IPv6 whilst using both protocols depending on whichever is more beneficial to the organization's needs. Thus instead of undergoing the costs for revamping an entire IPv4 network architecture firmly in place, the company has an option to invest within a comparatively lower budget and purchase a dual-stack router to service IPv4 and IPv6 clients in their internal network [6]. This deployment strategy thus provides both the performance and cost-effectiveness.

Thus despite the rapid exhaustion of IPv4 addresses by 2012, IPv4 could still remain a ubiquitous part of the world-wide-web. Organizations could still actively continue assigning IPv4 'private-addresses' for their internal networks and choose one of several migration strategies described above. Furthermore, the lack of economic incentive to deploy IPv6 [7] currently reinforces this theory that IPv4 remains to stay until at-least the very foreseeable future.

The remainder of this paper is organized as follows. Section II reviews literature on IPv4 and IPv6. Section III describes test bed and measurement procedure. The packet-generation and traffic-measuring mechanisms along with the evaluation

methodology is also discussed in this section. The experimental results and comparative analysis are presented in Section IV, and Section V concludes the paper.

II. LITERATURE REVIEW

In 2010, Chris Manford et al. [8] measured the performance of UDP protocol over both IPv4 and IPv6 using two client-server networks running Windows XP with Windows Server 2008 and Windows Vista with Windows Server 2008. Throughput and packet delay were measured using IP Traffic over a Fast Ethernet network. A total of 30 million packets were sent from source to destination to determine UDP performance. Their results revealed that Windows Vista-Server 2008 throughput to be between 26.9 to 85.5 Mbps for IPv4 and 26.3 to 84.1 Mbps for IPv6. For Windows XP-Server 2008, the throughput ranged between 27.9 to 85.5 Mbps for IPv4 and 24.5 to 85.2 Mbps for IPv6. They concluded that with the exception for the packet size of 384 bytes, there was negligible difference between the two operating systems. Windows Vista-Server 2008 however performed slightly better than Windows XP-Server 2008.

In 2010, Burjiz K. Soorty et al. [9] evaluated the performance of IPv4 and IPv6 over different cabling systems. Two commonly used industry cabling systems namely Category 5e and Category 6 were tested for IP performance over a Gigabit Ethernet LAN. Network throughput and packet delay were measured to evaluate network performance. Furthermore their study also took into account the additional parameter such as packet-length which could also affect the network performance. A detailed performance analysis was performed using different packet length (e.g. 128 to 1408 bytes). They found that UDP achieved higher throughput with Category 5e cabling than Category 6 thereby making Category 5e a much suitable preference for data traffic. Packet delay was found to be lower with Category 6 than Category 5e due to Category 6 providing about 12 dB (or 16 times) better Signal-to-Noise Ratio than Category 5e over a wide frequency range [9] thus making Category 6 a much viable preference for delay-sensitive traffic (e.g. VoIP).

In 2009, Chris Manford et al. [10] measured UDP performance over IPv4 and IPv6 on two peer-to-peer networks, namely Windows XP and Windows Vista. They focused on setting up a network commonly in small-to-medium sized businesses (SMB's). It was found that UDP throughput ranged from 26.5 Mbps to 85.6 Mbps for IPv4 and between 24.0 Mbps and 84.6 Mbps for IPv6 using Windows Vista. UDP throughput for Windows XP ranged from 21.8 Mbps to 86.0 Mbps for IPv4 and between 20.7 to 81.0 Mbps for IPv6. Overall, Windows Vista performs better than Windows XP and most corporate networks upgraded to Windows Vista for better performance.

In 2005, Tin-Yu Wu et al. [11] published a paper on the performance evaluation of IPv6 over Gigabit Ethernet. The main motivation was to analyze network efficiency under IPv6 structures. Unlike other studies, this paper compared the performance of IPv4 and IPv6 over inter-domain routing areas and included a Multi protocol label switching (MPLS)

backbone. They focused on core aspects of routing and switching performance rather than enterprise network OS performance. It was concluded that the throughput for UDP packets was 349 Mbps for IPv4 and 339 Mbps for IPv6. Throughput for TCP packets was 144 Mbps for IPv4 and 141 Mbps for IPv6.

Table 1 summarizes the key researchers and their main contributions in the performance evaluation of IPv6.

TABLE I. KEY RESEARCHERS AND THEIR MAIN CONTRIBUTIONS IN EVALUATING IPV6

Researchers	Year	Performance evaluation
C. Manford et al. [8]	2010	Measured network throughput and packet delay in Fast Ethernet for UDP using Windows XP-08 and Windows Vista-08.
B.K. Soorty et al. [9]	2010	Measured network throughput and packet delay in Gigabit Ethernet for TCP and UDP using Windows Vista.
C. Manford et al. [10]	2009	Measured network throughput and packet delay in Fast Ethernet for UDP using Windows XP and Windows Vista.
Tin-Yu Wu et al. [11]	2005	Measured network throughput and packet delay in Gigabit Ethernet for TCP and UDP using Fedora Core 2.0.

All the papers reviewed in this section considered only network throughput and packet delay. Performance metrics such as packet jitter and CPU utilization were not studied which might impact the performance of the two IP stacks according to an earlier study by Zeadally et al. [1].

Our main contribution in this paper is to obtain new results by investigating the drawbacks of implementing IPv6 with respect to QoS parameters packet delay, jitter, bandwidth and CPU utilization and further investigate which commonly used the latest operating systems would give the best performance. Using a real test bed we measure the QoS parameters to get an insight into the network performance using the latest Windows and Linux OSs.

III. TESTBED MEASUREMENT AND PROCEDURE

A. Testbed Configuration

In this research we study the UDP performance of IPv6 using the latest Windows and Linux systems. The network topology is a peer-to-peer Gigabit Ethernet client and server machines. No routers, switches or hubs were used in the experimental setup so as to ensure that there was no latency experienced on the network. Furthermore all services (running on default) consuming network bandwidth and/or CPU resources were disabled to get unbiased and more accurate results. No third-party applications were used to optimize or influence network performance in any way.

Each workstation was separated from the other by a distance of approximately one meter. This was done so as to maintain consistency with earlier research [8-10] and thus produce results indicative for a fair comparison of the same. The client and server machines were connected using a Category 6

Crossover UTP (Unshielded Twisted Pair) cable maintaining EIA/TIA 568-B wiring configuration (Fig. 1).

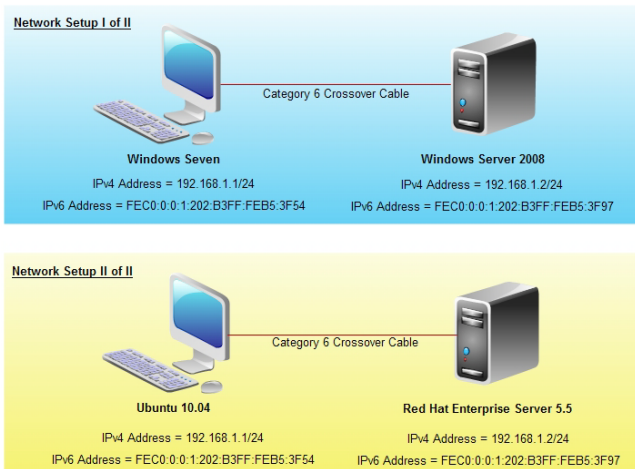


Figure 1. Network Testbed.

The hardware benchmark consisted of four workstations, all of which surpassed the minimum and recommended settings for the applicable OSs tested on them. Workstations 1 and 2 were each running Intel® Core™ 2 Duo processors with 4 GB 800 MHz DDR-2 Corsair® RAM modules for Windows 7 and Windows Server 2008. Workstation three was running an Intel Core 2 Duo processor with 2 GB 800 MHz DDR-2 RAM modules for Ubuntu 10.05. Workstation 4 was a Lenovo T40 laptop that was used for Red Hat Server 5.5 after repeated tests revealed that the native hardware configuration met the satisfied recommended OS settings. Workstations were equipped with the Realtek PCIe GBE Family Controller NIC (Network Interface Controller) card to carry out the Gigabit Ethernet evaluations on the network.

B. Measurement Tools and Metrics

Several data-generating and traffic-measuring tools were researched for the purpose of evaluating IPv6 on Windows and Linux operating systems.

For Windows, IP Traffic [12] was selected as the preferred tool of choice due to its extensive history as a widely-used tool and for its overall accuracy in evaluating network performance. IP Traffic is also the only program to work on Windows 7 and Windows Server 2008. Furthermore, several publications such as [8-10] used IP Traffic to evaluate performance of IPv4 and IPv6. It has also been commonly used on wireless LAN evaluations [13, 14].

For Linux systems, a heavily modified tool of Iperf [15] was used to evaluate performance of IPv4 and IPv6. Iperf is an open-source network performance measurement tool that can create TCP and UDP data streams and measure the throughput and delay of a network that is carrying them. Iperf was used in [16] as the primary evaluation tool for IPv6 performance.

For each evaluation-run performed, IP Traffic sent a total of one million packets. Ten such runs were recorded per protocol for each Windows-based operating-system. A total of ten million packets were thus sent before each protocol's throughput and delay was recorded for every packet-size using a particular Windows based operating-system. A similar approach was adapted for recording performance of each Linux-based operating system using Iperf. This was done in-order to maintain accuracy and consistency of results. A standard deviation of less than 10% was implemented to accurately measure overall network performance.

IV. EXPERIMENTAL RESULTS

Figure 2 compares UDP throughput of IPv4 (2A) and IPv6 (2B) using Windows and Linux. We observe that the highest throughput is achieved using Ubuntu 10.04-Red Hat Server 5.5 client-server network than Windows 7-Server 2008. For example, for IPv4 using Ubuntu 10.04-Red Hat Server 5.5 the throughput is 774.46 Mbps (6.53% increase from Windows 7-Server 2008.) at packet length of 1408 bytes.

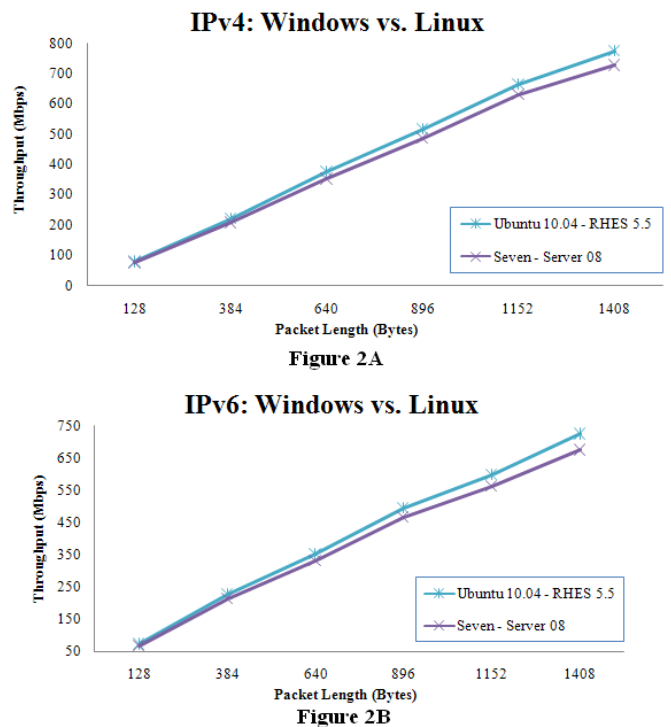


Figure 2: Throughput comparison of Windows and Linux client-server operating systems on IPv4 (2A) and IPv6 (2B)

For IPv6 using Ubuntu 10.04-Red Hat Server 5.5, the highest throughput is 725.00 Mbps (7.48% increase from Windows 7-Server 2008) at packet-length of 1408 bytes.

As shown in Fig. 3 and Table 2, UDP throughput is higher on IPv4 than IPv6 using both Windows and Linux systems. This is due to the IPv6 header being six times larger than the IPv4 header combined with the fact that the more efficient and

simplified header of IPv6 does not have significant play with the connectionless nature of the UDP protocol.

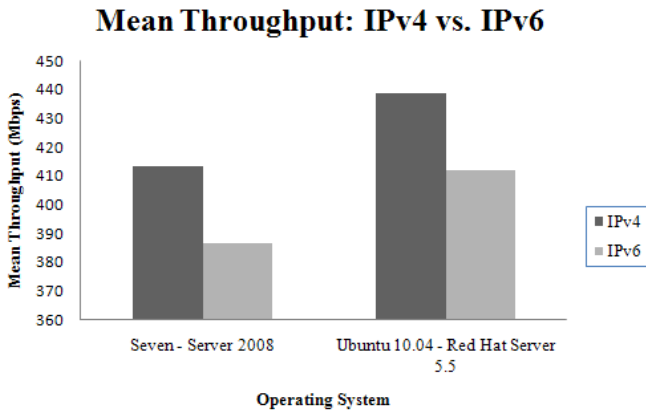


Figure 3: Average Throughput Comparison of IPv4 and IPv6 on Windows and Linux client-server operating systems

Figure 4 compares the UDP packet delay between each client-server operating system on IPv4 and IPv6. The lowest packet delay was recorded at 1.27 ms for the smallest packet-size of 128 bytes on Windows 7-Server 2008 compared to 1.37 ms for Ubuntu 10.04-RHES 5.5 at the smallest packet length of 128 bytes.

The highest packet delay rate was recorded at 4.92 ms for the largest packet length of 1408 bytes on Ubuntu 10.04-RHES 5.5 compared to 4.67 ms for Seven-Server 2008 also recording at the largest packet-size of 128 bytes.

Despite Ubuntu 10.04-Red Hat Server 5.5 producing a significantly higher throughput for UDP than using Windows 7-Server 2008, the difference in UDP delay between the two networks is considerably low. This is probably due to Ubuntu having a more efficient socket layer that responds faster to kernel switches during system calls.

TABLE 2: IPv4/IPv6 THROUGHPUT COMPARISON OF WINDOWS AND LINUX

Operating Systems	Throughput (Mbps)	
	IPv4	IPv6
Ubuntu 10.04 - Red Hat Server 5.5	438.68	411.83
Windows 7 Server 2008	412.92	386.33

The second metric measured and evaluated in this paper was packet delay. UDP packets were sent across the IP networks and round trip time was measured in milliseconds.

Figure 5 compares packet delay for IPv4 and IPv6 over each client-server network. On Windows 7 Server 2008, the lowest packet delay is 1.27 ms for both IPv4 and IPv6. The highest packet delays are 4.67 ms and 4.09 ms, for IPv4 for IPv6 respectively. For example, On Ubuntu 10.04-RHES 5.5, the lowest RTT is 1.37 ms for IPv4 and 1.35 ms for IPv6. The highest delay is 4.92 ms for IPv4 and 4.33 ms for IPv6.

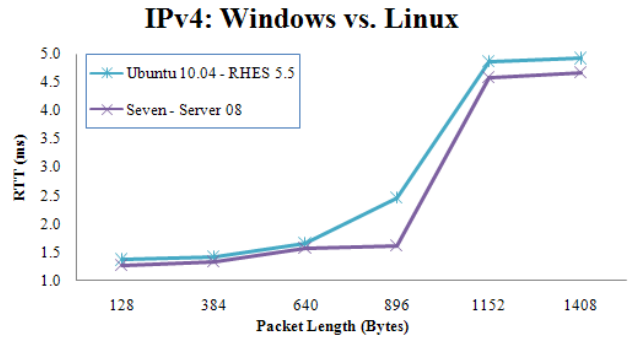


Figure 4A

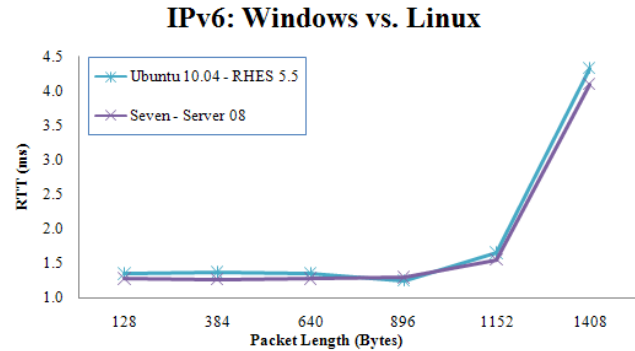


Figure 4B

Figure 4: Packet delay comparison of Windows and Linux client-server operating systems on IPv4 (4A) and IPv6 (4B)




Figure 5: Mean packet delay comparison of IPv4 and IPv6 on Windows and Linux client-server OSs.

We also observe that the lowest packet delay is achieved for shorter packet length. This is mainly due to throughput being comparatively lower on smaller packets. With UDP having no error-correction mechanism there is no significant overhead in relation to the payload and therefore no restriction in throughput from reaching its maximum. This results in relatively lower delay compared to larger packets. As shown in Fig. 5 and Table 3, the packet delay is significantly lower for IPv6 than IPv4 for all three networks.

TABLE 3: IPV4/IPV6 DELAY COMPARISON ON WINDOWS AND LINUX

Operating Systems	Packet delay (ms)	
	IPv4	IPv6
Ubuntu 10.04 - Red Hat Server 5.5	2.78	1.88
Windows 7 - Server 2008	2.51	1.79
Vista-Server 2008 [17]	2.33	1.58

The third metric measured and evaluated in this paper was packet jitter. UDP packets were sent across the IP networks and packet jitter was measured at the receiving node (server) for each network scenario.

Figure 6 compares UDP packet jitter for IPv4 and IPv6 over the two client-server networks. As shown in Fig. 6, packet jitter is lower on Windows 7-Server 2008 for both the IPv4 and IPv6.




Figure 6: Jitter comparison of Windows and Linux client-server operating systems for IPv4 and IPv6

The fourth metric measured and evaluated in this paper was CPU utilization for transferring data across the network (client host). Figure 7 shows CPU usage for transferring UDP packets. We observe that CPU utilization is slightly lower with Ubuntu 10.04 than Windows 7 for IPv4 and IPv6. For example, minimum utilization for Ubuntu is 23.80% for IPv4 and 19.97% for IPv6 compared to 27.66% for IPv4 and 25.34% for IPv6 using Windows 7.

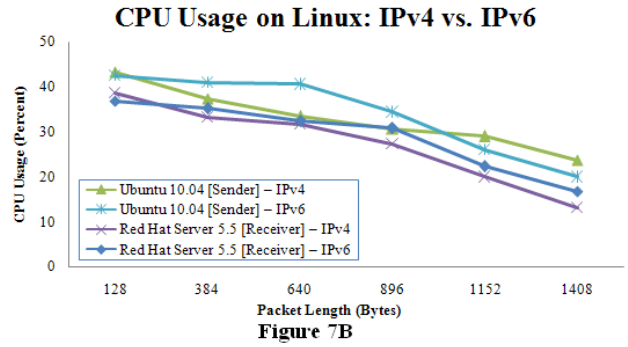



Figure 7: CPU Usage comparison of Windows (7A) and Linux (7B) client-server operating systems on IPv4 and IPv6

V. CONCLUSIONS AND DISCUSSION

In this paper we have studied UDP performance of IPv6 in a peer-to-peer Gigabit Ethernet LAN using the latest MS Windows and Linux operating systems. Our findings based on empirical study are summarized below.

- UDP throughput for IPv6 was highest on the Linux network running Ubuntu 10.04 with Red Hat Server 5.5 (Ubuntu 10.04-Red Hat Server 5.5 offers an average throughput of 411.83 Mbps compared to 386.33 Mbps on Windows 7-Server 2008 network).
- UDP throughput for IPv4 was also highest on the Linux network running Ubuntu 10.04 with Red Hat Server 5.5 (Ubuntu 10.04-Red Hat Server 5.5 offers an average throughput of 438.68 Mbps compared to 412.92 Mbps on the Windows 7 Server 2008).
- UDP mean packet delay was lowest on Windows 7-Server 2008 for IPv4 (2.51 ms compared to 2.78 ms for Ubuntu 10.04-Red Hat Server 5.5). For IPv6, packet delay was also lowest on Windows 7-Server 2008 (1.79 ms compared to 1.88 ms for Ubuntu 10.04-Red Hat Server 5.5).
- UDP packet jitter was lowest on Windows 7-Server 2008 than Ubuntu 10.04-Red Hat Server 5.5.
- CPU utilization to transfer UDP packets across the network was more efficient over Ubuntu 10.04 compared to Windows 7-Server 2008.

In conclusion, we found that Linux network running Ubuntu 10.04 with Red Hat Server 5.5 performed significantly better on IPv6 than the Windows network running Windows 7 with Windows Server 2008. Even though packet delay and packet jitter were comparatively higher on Ubuntu 10.04-Red Hat Server 5.5, the overall performance was better on Linux due to its significant throughput gain than Windows 7.

REFERENCES

- [1] S. Zeadally and L. Raicu, "Evaluating IPv6 on Windows and Solaris," *IEEE Internet Computing*, vol. 7, no. 3, pp. 51-57, 2003.
- [2] A. A. Mahimkar, H. H. Song, Z. Ge, A. Shaikh, J. Wang, J. Yates, Y. Zhang, and J. Emmons, "Detecting the performance impact of upgrades in large operational networks," *SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 4, pp. 303-314, 2010.
- [3] M. Mackay and C. Edwards, "A Managed IPv6 Transitioning Architecture for Large Network Deployments," *IEEE Internet Computing*, vol. 13, no. 4, pp. 42-51, 2009.
- [4] P. Savola, "Observations of IPv6 traffic on a 6to4 relay," *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 1, pp. 23-28, 2005.
- [5] L. Sang-Do, S. Myung-Ki, and K. Hyoung-Jun, "The implementation of ISATAP router," presented at the 8th International Conference on Advanced Communication Technology, Phoenix Park, Gangwon-Do, Korea, February 20-22, 2006, pp. 3 pp.-1163.
- [6] U. Narayanan and X. Jiang, "Signaling Cost Analysis of Handoffs in a Mixed IPv4/IPv6 Mobile Environment," presented at IEEE Global Telecommunications Conference (GLOBECOM '07), Washington, DC, November 26-30, 2007, pp. 1792-1796.
- [7] M. Kuhne, "The Seven Stages of IPv6 Adoption," *IETF Journal*, vol. 5, no. 1, pp. 14-17, 2009.
- [8] C. Manford, S. S. Kolahi, B. K. Soorty, Z. Qu, and N. Chand, "UDP Performance Comparison of IPv4 and IPv6 using Windows Vista-Server2008 and Windows XP-Server2008 LANs," presented at Advances in Computer Science and Engineering (ACSE 2010), Sharm El Sheikh, Egypt, March 15-17, 2010, pp.
- [9] B. K. Soorty, S. S. Kolahi, N. Chand, and Q. Zhang, "Performance Comparison of Category 5e vs. Category 6 Cabling Systems for both IPv4 and IPv6 in Gigabit Ethernet," presented at the IEEE 10th International Conference on Computer and Information Technology, Bradford, UK, June 29-July 1, 2010, pp. 1525-1529.
- [10] C. Manford, S. S. Kolahi, B. K. Soorty, Q. Zhang, and N. Chand, "UDP Performance Analysis of IPv4 and IPv6 on Windows Vista and Windows XP over Fast Ethernet Peer-Peer LAN," presented at the 4th International Conference on Ubiquitous Information Technologies & Applications, Fukuoka, Japan, December 20-22, 2009, pp. 1-4.
- [11] T.-Y. Wu, H.-C. Chao, T.-G. Tsuei, and Y.-F. Li, "A measurement study of network efficiency for TWAREN IPv6 backbone," *International Journal of Network Management*, vol. 15, no. 6, pp. 411-419, 2005.
- [12] ZTI-Telecom. (2011). *IP Traffic - Test & Measure*, from http://www.zti-telecom.com/EN/IPTraffic_TM_KeyFeatures.html
- [13] N. Baghaei and R. Hunt, "IEEE 802.11 wireless LAN security performance using multiple clients," presented at the 12th IEEE International Conference on Networks Singapore, November 16-19, 2004, pp. 299-303.
- [14] S. S. Kolahi, S. Narayan, D. D. T. Nguyen, Y. Sunarto, and P. Mani, "The impact of wireless LAN security on performance of different Windows operating systems," presented at IEEE Symposium on Computers and Communications, Marrakech, Morocco, July 6-9, 2008, pp. 260-264.
- [15] Jdugan and Mitchkutzko. *Iperf*, 2011, from <http://sourceforge.net/projects/iperf/>
- [16] S. Narayan, S. S. Kolahi, Y. Sunarto, D. Nguyen, and P. Mani, "Performance comparison of IPv4 and IPv6 on various windows operating systems," presented at the 11th International Conference on Computer and Information Technology Khulna, Bangladesh, December 24-27, 2008, pp. 663-668.
- [17] S. S. Kolahi, N. Chand, B. K. Soorty, and Q. Zhang, "Performance Comparison of IPv4 and IPv6 in Peer-Peer and Client Server Local Area Networks," presented at the IEEE 10th International Conference on Computer and Information Technology, Bradford, U.K., June 29-July 1, 2010, pp. 41-45.