Implementation of IPv6 ToS over ATM Network

Ricky Ng, Danny Yip, and Ljiljana Trajkovic
School of Engineering Science
Simon Fraser University
Vancouver, British Columbia
Canada V5A 1S6
E-mail: {rickyn, ymyip, ljilja}@cs.sfu.ca

Abstract

The current Internet infrastructure uses Internet Protocol version 4 (IPv4) that only supports the "best-effort" Quality-of-Service (QoS). While this simplification of QoS requires small processing latency in routers, there is no distinction between packets with various delay requirements. The introduction of the Type of Service (ToS) field in Internet Protocol version 6 (IPv6), addresses this deficiency. The feasibility of mapping IPv6 ToS onto various Asynchronous Transfer Mode (ATM) service categories has been recently suggested. We investigate the application and the mapping of the IPv6 ToS onto various ATM QoS classes. Our goal is to use the OPNET simulation tool to verify the interoperability of IPv6 ToS and ATM QoS service categories. In our implementation, IPv6 packets may be switched and delivered to their destination based on the specified ToS header fields. The simulation results indicate that such an IPv6-over-ATM network could deliver IPv6 packets according to their specified ToS, and that packets with distinct ToS will experience different end-to-end network delays.

1 Introduction

IPv6 [1, 2, 3] is a new version of the Internet Protocol (IP), designed as a successor to IPv4. It was intended to be evolutionary different from IPv4, rather than to quickly replace IPv4. Since IPv6 was designed to enable a smooth transition from IPv4, many IPv4 functions were kept in IPv6. However, there are several IPv4 functions that were eliminated. The changes from IPv4 to IPv6 fall into the following categories:

- *Expanded Routing and Addressing Capabilities*:
  Since the number of computers connected to the Internet increases exponentially, IPv4's 32-bit wide IP addresses are used up quickly and can no longer meet user demands. As a result, IPv6 was designed to have 128 bits (instead of 32 bits) addressing space in order to support additional levels of addressing hierarchy and a much greater number of addressable nodes.
- *Header Format Simplification*:
  In addition to the expanded addressing capabilities, IPv6's header format was also designed to minimize the processing cost of packet handling and to keep the bandwidth of the IPv6's overhead as low as possible. Even though the IP address space in IPv6 has increased, the IPv6 header is only twice the size of the IPv4 header because certain IPv4 header fields were omitted and made optional.
- *Improved Support for Options*:
  IPv6 also improves the way IP header options are encoded. Such improvement allows for more efficient forwarding, less stringent limits on the length of options, and greater flexibility for introducing new future options.
- *Authentication and Privacy Capabilities*:
  One of the major improvements of IPv6 is the introduction of authentication, data integrity, and confidentiality. They are considered to be basic elements of IPv6, included in all implementations. In the future, IPv6 packets with high confidentiality will only be transmitted through routers that are specified by the sender.
- *Quality-of-Service Capabilities*:
  Lastly, IPv6 has the capability to enable the labeling of packets that belong to particular traffic "flows." IPv6 header has a field called Type of Service (ToS) that allows the sender to label the packets sent into the network. The IPv6 network handles the packets according to the specified label and delivers certain quality of service to the sender. Although similar capability was implemented in IPv4 (using a header field called "service type"), it was not well defined at the time when the protocol was widely adopted and implemented. As a result, today's networking equipment simply ignores this field and all IPv4 packets are treated equally.

The goal of our project is to investigate different ToS values offered by IPv6 and to examine how to take advantage of the service categories and Quality of Service (QoS) offered by Asynchronous Transfer Mode (ATM) networks [4, 5], so that IPv6 packets can be switched and delivered to the destination based on their specified ToS header fields [6]. Such an IPv6-over-ATM network could deliver IPv6 packets according to their specified ToS, and packets with distinct ToS will experience different end-to-end network delays.

In our implementation, we modified three existing OPNET node models [7]: *atm_uni_src*, *atm_crossconnect*, and *atm_uni_dest*. They are used to model an ATM network that allows the ATM sources (*atm_uni_src*) to generate raw packets employing any of the five ATM service categories. These packets, encapsulated by the ATM Adaptation Layer 5 (AAL5), are segmented into ATM cells (*atm_crossconnect*) and are sent to an ATM switch that routes them to the destination (*atm_uni_dest*) where they are re-assembled into AAL5 packets. We implemented an IPv6-over-ATM interface module that determines the ATM Switched Virtual Circuit (SVC) that each IPv6 packet should occupy when being delivered to its destination. The simulated network consists of a host that uses the new module. Simulation results revealed that IPv6 packets with higher priority (specified by the ToS field) experience smaller end-to-end packet delay. These results indicate the feasibility of using the ToS header byte in IPv6 packets in meeting various ATM QoS requirements.

2 Mapping IPv6's ToS to ATM Service Category

The 4-bit wide ToS field in the IPv6 header enables a source to identify the desired delivery priority of each packet relative to other packets from the same source. According to the IPv6 Specification RFC1883 [3], the ToS field is further separated into two categories: 1) Congestion-controlled traffic and 2) Non-congestion-controlled traffic.

Congestion-controlled traffic

Congestion-controlled traffic refers to packets for which the source "backs off" in response to congestion. Such mechanism exists in Transmission Control Protocol (TCP). Because of the nature of congestion-controlled traffic, a variable delay in the delivery of packets is acceptable. IPv6 defines the following categories of congestion-controlled traffic:

- *ToS = 0, Uncharacterized traffic*:
  If the upper-layer application gives IPv6 no guidance regarding traffic priority, then those packets will be assigned the lowest-priority value and will be delivered based on the available bandwidth.
- *ToS = 1, Filler traffic*:
  These are applications that generate traffic handled in the background, while other types of traffic are delivered. Those applications include newsgroups messages and electronic mail.
- *ToS = 2, Unattended data transfer*:
  These are applications that generate packets that are not delivered instantaneously. The best example of this category is electronic mail. Generally, the user does not wait for the transfer to be completed, and, therefore, those packets can experience longer delay.
- *ToS = 4, Attended bulk transfer*:
  These are applications that may involve the transfer of a large amount of data. Those applications include File Transfer Protocol (FTP) and Hypertext Transfer Protocol (HTTP). During these application sessions, users are generally prepared to accept a longer delay, and, therefore, packets with ToS set to 4 can be delivered at a lower priority than packets with ToS equal to 6 and higher.
- *ToS = 6, Interactive traffic*:
  After Internet control traffic, the most important traffic to support is interactive traffic, such as telnet. Rapid response time during interactive session is crucial because a user is interacting with the host in real-time, and, therefore, delay should be minimized.
- *ToS = 7, Internet control traffic*:
  This is the most important traffic to deliver, especially during times of high congestion. For example, protocols such as Open Shortest Path First (OSPF) and Simple Network Management Protocol (SNMP) need to receive updates concerning traffic conditions and report congestion to other network management application in order to adjust routing path and perform dynamic reconfiguration to relieve congestion conditions.

Table 1 summarizes the different ToS values for congestion-controlled traffic and examples of applications that apply to each ToS category. According to the IPv6 Specification RFC1883 [3], ToS values of 3 and 5 are reserved. Therefore, source will not generate packets with ToS set to these values.

ToS	Description
0	Uncharacterized traffic
1	"Filler" traffic (netnews)
2	Unattended data transfer (email)
3	Reserved
4	Attended bulk transfer (FTP, NFS)
5	Reserved
6	Interactive traffic (telnet)
7	Internet control traffic (routing protocols, SNMP)

Table 1: ToS description for congestion-controlled traffic.

Non-congestion-controlled traffic

Non-congestion-controlled traffic requires relatively constant delivery delay. Examples are real-time video and audio over User Datagram Protocol (UDP), which maintain smooth delivery flow and do not require retransmission because it is not useful to re-transmit voice or real-time video frames that were previously dropped and replay them in the middle of a real-time transmission.

Eight levels of priority are allocated for this type of traffic from the lowest priority 8 (most willing to discard) to the highest priority 15 (least willing to discard). In general, the criterion is how much the quality of the received traffic the user can tolerate when packets are dropped. For example, a telephone voice conversation would typically be assigned a high priority because humans are more sensitive to audio signal loss. On the other hand, video signal contains redundant information between frames, and the loss of a few packets will probably not be noticeable; therefore, this traffic is assigned a relatively low priority. Table 2 summarizes the different ToS values for non-congestion-controlled traffic and examples of applications that apply to each ToS categories.

Although there is a distinction between the two ToS traffic categories, there is no priority relationship implied between the congestion-controlled priorities and the non-congestion-controlled priorities. Priorities are relative only within each category.

ToS	Description
8	Lowest priority, most willing to discard (high fidelity video traffic)
9-14	…
15	Highest priority, least willing to discard (low fidelity video traffic)

Table 2: ToS description for non-congestion-controlled traffic.

3 IPv6's ToS vs. ATM's QoS Categories

This subsection discusses different options for selecting ATM's service category to carry different IPv6's ToS traffic. ATM offers five different service categories:

- *Constant Bit Rate (CBR)*
- *Real-Time Variable Bit Rate (RT-VBR)*
- *Non-Real-Time Variable Bit Rate (NRT-VBR)*
- *Available Bit Rate (ABR)*
- *Unspecified Bit Rate (UBR)*.

Based on the characteristics of each service category and the suggestions that are consolidated from different sources, the following summarizes one possible mapping between IPv6's ToS and ATM's service categories:

Congestion-Controlled Traffic (ToS = 0 to 7) over ABR or UBR

The use of ABR or UBR to carry congestion-controlled traffic was widely suggested because both ABR and UBR exhibit the "best-effort" behavior that IP provides. Using the ABR service to carry congestion-controlled traffic was relatively less recommended because ABR is also embedded within congestion-controlled mechanism. If congestion-controlled traffic is carried over a connection that also has congestion-controlled mechanism, the source will not be notified right away when there is congestion in the connection. During congestion, ABR's congestion-controlled mechanism would attempt to relieve the congestion condition, which delays upper layer congestion-controlled applications such as TCP from slowing down their transmission. However, the use of ABR can minimize the loss in the ATM network (due to ABR's congestion-controlled mechanism) and therefore maximizes the throughput and network utilization. UBR is highly suggested to carry IPv6 packets generated from applications that can tolerate higher delay and jitter, which are the characteristics of those packets with ToS equal to 0 to 7.

Non-congestion-controlled traffic (ToS = 8 to 15) over NRT-VBR, RT-VBR, or CBR

The CBR, RT-VBR, and NRT-VBR service categories are suggested to carry voice and video packets. According to the IPv6 specification, those packets would have their ToS header fields set to values between 8 and 15 (inclusive). The rationale behind the selection is that CBR service can carry all the data through the ATM network with minimal loss at a constant rate, which is the characteristic of most real-time voice and video signals. However, the use of CBR does not allow access to unused ATM network bandwidth on an as-needed basis and therefore does not maximize network utilization. In the case of RT-VBR and NRT-VBR, an average rate and burstiness of the traffic arrival are specified, which are ideal to carry bursty traffic such as compressed-voice and compressed-video signals, where packet's size differ from packet to packet. The use of RT-VBR and NRT-VBR also allows access to excess bandwidth by using the Cell Loss Priority (CLP) bit in the ATM header, and, therefore, achieves better network utilization.

Table 3 summarizes what we believe is the most appropriate mapping between IPv6's ToS and ATM's service categories.

	ToS	Description	Suggested service category
Congestion-controlled traffic	0	Uncharacterized traffic	UBR/ABR
	1	"Filler" traffic (netnews)	UBR/ABR
	2	Unattended data transfer (email)	UBR/ABR
	3	Reserved	UBR/ABR
	4	Attended bulk transfer (FTP, NFS)	UBR/ABR
	5	Reserved	UBR/ABR
	6	Interactive traffic (telnet)	UBR/ABR
	7	Internet control traffic (routing protocols, SNMP	UBR/ABR
Non-congestion-controlled traffic	8	Lowest priority, most willing to discard (high fidelity video traffic)	NRT-VBR
	9-14	…	NRT-VBR/ RT-VBR/ CBR
	15	Highest priority, least willing to discard (low fidelity video traffic)	CBR

Table 3: Suggested mapping between IPv6 ToS and ATM service categories.

4 Implementation of IPv6-over-ATM OPNET Model

The OPNET node models that we used to implement IPv6's ToS over different ATM service categories are:

- *atm_uni_src*
- *atm_crossconnect*
- *atm_uni_dest*.

These models have already been implemented in ATM network that allows the source (*atm_uni_src*) to generate raw packets over a unique Switched Virtual Circuit (SVC) or Permanent Virtual Circuit (PVC) with any of the five ATM service categories. The packets generated by the source will then be encapsulated into ATM Adaptation Layer 5 (AAL5) packets and get segmented into ATM cells before being sent to the ATM switch (*atm_crossconnect*). The ATM switch will then switch the received ATM cells to the destination (*atm_uni_dest*), which will re-assemble the received ATM cells into AAL5 packets. Finally, the raw packets will be extracted from the AAL5 packets.

The following subsections describe the changes made to the above OPNET models in order to implement an IPv6-over-ATM network that will switch IPv6 packets based on their specified ToS.

IPv6 Source

Originally, the *atm_uni_src* can only set up one SVC per node and the node can only generate raw packets. Figure 1 shows the part of the *atm_uni_src* node model that was modified.

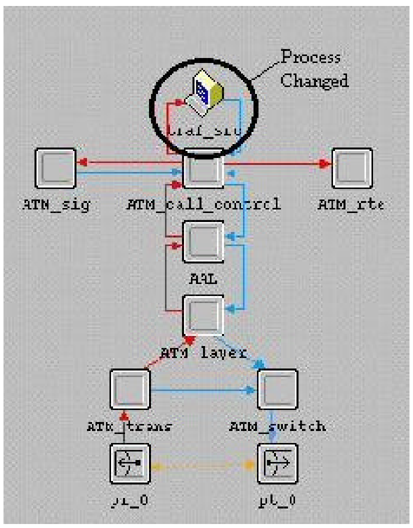

Figure 1: Modified *atm_uni_src* node model.

The process, "traf_src", was modified such that the source can now generate IPv6 packets with their ToS header fields uniformly assigned between 0 to 15, except for ToS = 3 and ToS = 5 that are reserved. Figure 2 shows the IPv6 packet that is created. The packet generator that creates packets with various ToS uses this packet format.

Figure 2: IPv6 header format.

Furthermore, by changing the "traf_src" process, the node can now create multiple SVC's with different service categories instead of the original implementation that could support only one SVC per node. This modification is required in order to allow the packet generator to transmit IPv6 packets over multiple SVC's within the same node. A new user attribute called "IPv6 ToS Mapping" is also added to the source. This new attribute allows the user to select the ATM service category to map IPv6's ToS. At simulation time, depending on the user's selection, SVC's of the specified ATM service category will be set up before packets are generated. For example, if a user selects to map all congestion-controlled traffic (ToS = 0-7) to UBR and all non-congestion-controlled traffic (ToS = 8-15) to RT-VBR, two SVC's will be set up at simulation time, one for UBR and the other one for RT-VBR. Once all SVC connections are established, IPv6 packets with different ToS will be generated and transmitted to the corresponding SVC according to the user's selection. Figures 3 to 6 show the modified node attribute window that contains the new user attribute "IPv6 ToS Mapping" and its subsequent windows, which allow the user to select the ToS mapping to various ATM service categories.

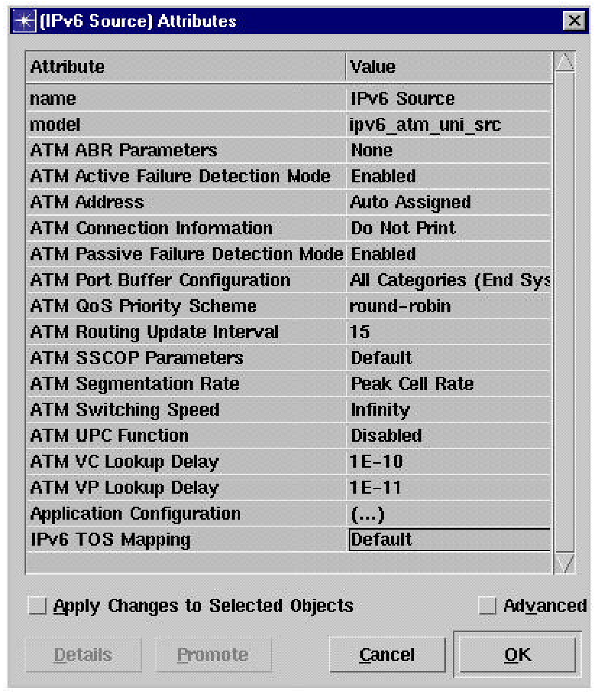

Figure 3: IPv6 ToS mapping attribute.

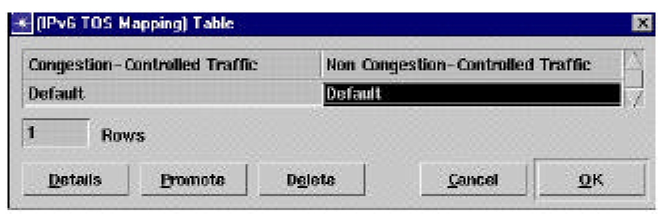

Figure 4: Choice of two different traffic categories.

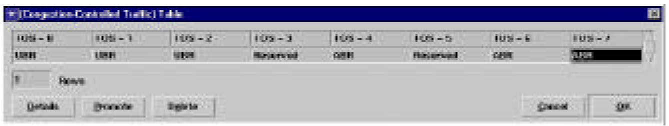

Figure 5: ATM service category mapping for congestion-controlled traffic.

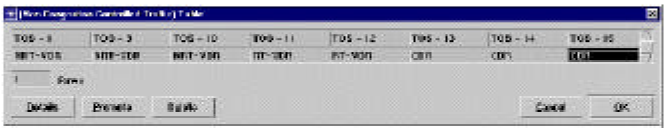

Figure 6: ATM service category mapping for non-congestion-controlled traffic.

By adding this new user-attribute to the IPv6 source node, users can experiment the effects and consequences of sending IPv6 packets with different ToS over SVC with different service categories.

IPv6-over-ATM Switch

The *atm_crossconnect* node is another ATM node model from OPNET standard library that is enhanced. Figure 7 shows the modified switch module in the *atm_crossconnect* node.

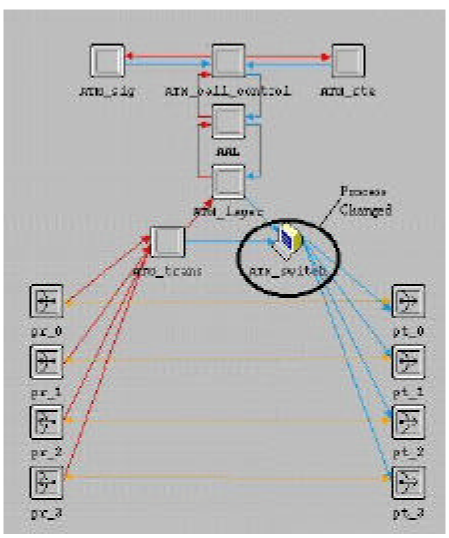

Figure 7: Modified *atm_crossconnect* node model.

The "ATM_switch" module was modified because the weighted round robin algorithm, which resides in the "ATM switch" module, did not discriminate in delivering cells of different service categories under non-congested condition.

The ATM "dequeue" function inside the "ATM_switch" module was modified in order for queues to be serviced according to the weights (minimum guaranteed bandwidth) that are assigned by the user under any traffic conditions. Figures 8 and 9 show the "ATM switch" attribute windows that allow configuration of the weight of the queue (minimum guaranteed bandwidth) by the user.

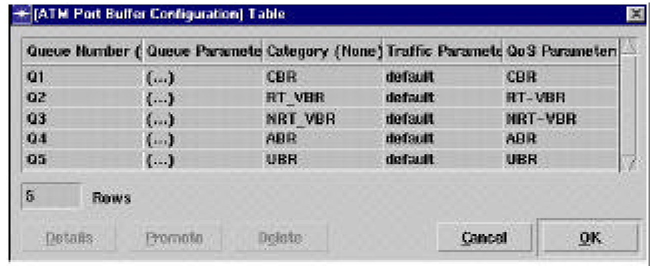

Figure 8: ATM port buffer configuration attribute.

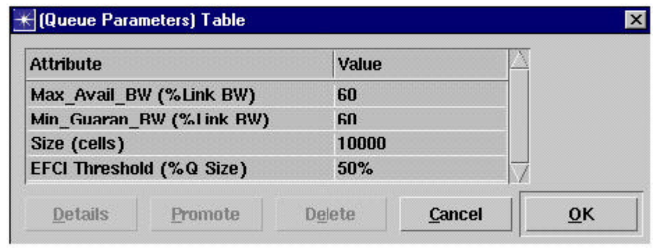

Figure 9: Queue parameters of a service category queue.

The "dequeue" function is changed so that it will first determine which queue has been assigned the largest "minimum guaranteed bandwidth." The remaining queues are assigned a relative "wait-counter" based on the "maximum" value. For example, let the CBR queue have the largest minimum guaranteed bandwidth and be assigned 60% of the bandwidth, while RT-VBR and UBR queue are assigned 30% and 10%, respectively. The relative "wait-counter" of the CBR queue is 1, and the relative "wait-counter" of the RT-VBR and UBR queue are 2 and 6, respectively. A cell in the queue will not be serviced until the "wait-counter" weight of the queue has been consumed. Hence, a CBR cell will be serviced in each of its designated round, while RT-VBR will be serviced every 2 times, and UBR will be serviced every 6 times. Although the dequeuing algorithm is not very efficient, it implements a weighted round robin QoS priority algorithm to allow the switch to discriminate ATM cells based on the minimum guaranteed bandwidth of the corresponding service queue under non-congested condition.

IPv6 Destination

The *atm_uni_dest* node model from the ATM OPNET library is also modified in order to capture the end-to-end packet delay experienced by IPv6 packets with different ToS field. Figure 10 shows the modified module in the *atm_uni_dest* node. The "traf_sink" process is responsible for processing the received packet and for updating the statistics.

The "traf_sink" process is enhanced so that it will first examine the ToS field when an IPv6 packet is received, then calculate the end-to-end delay and finally log the result to the corresponding statistics handle. The interface of the process is also modified to allow user to select the end-to-end packet delay statistics that he/she wants to collect before running the simulation. Figure 11 shows the "Collect Individual Statistics" window of the sink node with the enhanced statistics handles.

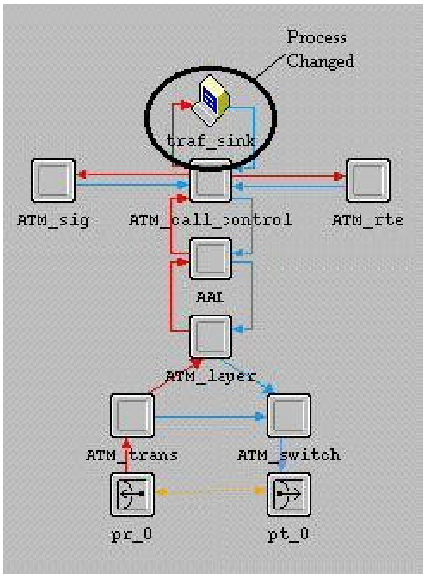

Figure 10: Modified *atm_uni_dest* node model.

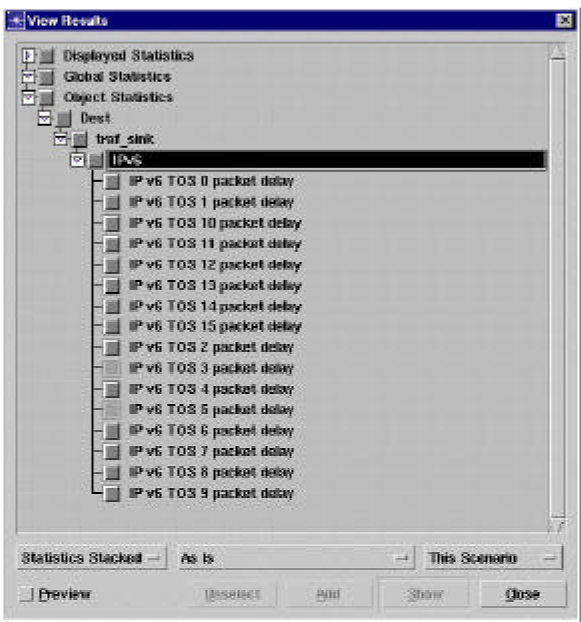

Figure 11: New IPv6 statistics handles at the sink.

5 Simulation Results

Figure 12 shows the network topology of the IPv6-over-ATM network simulation. Table 4 summarizes the network settings selected for the simulation.

The mapping of the ToS value to the ATM service categories for this simulation scenario is indicated in Table 5. The traffic contract and QoS parameters of different ATM service categories are shown in Tables 6 and 7, respectively. In Table 6, ATM traffic contract parameters are Sustainable Cell Rate (SCR), Peak Cell Rate (PCR), and Maximum Burst Size (MBS). In Table 7, QoS parameters are peak-to-peak Cell Delay Variation (ppCDV), maximum Cell Transfer Delay (maxCTD), and Cell Loss Ratio (CLR). Finally, Table 8 shows the percentage of bandwidth assigned to each service category queue in the IP-over-ATM switch node that will determine the relative "wait-counter" of each ATM service category.

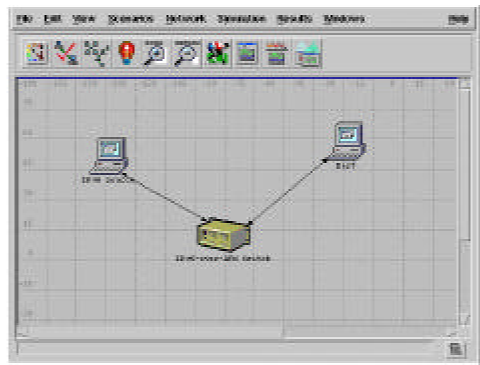

Figure 12: IPv6-over-ATM network topology.

Arrival rate (source)	0.5 Mbps
Arrival distribution (source)	Exponential
Transmission size (source)	3 Mbytes
Packet size	125 bytes
Simulation time	2 minutes (Simulation will end when either the simulation time expires or when the specified transmission size has been delivered.)

Table 4: Simulation scenario.

ToS	CBR	RT-VBR	NRT-VBR	ABR/UBR
0-7				X
8-10			X	
11-13		X		
14-15	X			

Table 5: Mapping of ToS to ATM service category (configurable).

Service category	SCR	PCR	MBS
CBR	0.5 Mbps	1 Mbps	10 cells
RT-VBR	0.5 Mbps	1 Mbps	10 cells
NRT-VBR	0.5 Mbps	1 Mbps	10 cells
ABR	0.5 Mbps	1 Mbps	10 cells
UBR	0.5 Mbps	1 Mbps	10 cells

Table 6: ATM traffic contract (non-configurable). Values are derived from the OPNET default model.

Figure 13 indicates the packet end-to-end delay for IPv6 packets with three ToS values: 0, 8, and 12. The collected statistics show that IPv6 packets that are carried over an ATM link with larger weight (% of the bandwidth) will experience shorter end-to-end delay. One exception is the ABR traffic that is outperformed by the UBR traffic when both queues are assigned the same weight. (The reason of such discrepancy may be due to the flow control algorithm employed by the ABR traffic, which introduces additional overhead.) These results indicate the feasibility of assigning higher priority IPv6 packets onto higher QoS ATM SVC's in guaranteeing a bounded delay that is a critical parameter for real time voice and video applications.

Service category	ppCDV	maxCTD	CLR
CBR	3 msec	400 msec	3E-07
RT-VBR	10E+06 msec	10E+06 msec	3E-07
NRT-VBR	Unspecified	Unspecified	1E-05
ABR	10E+06 msec	10E+06 msec	1E-05
UBR	Unspecified	Unspecified	Unspecified

Table 7: ATM quality of services (non-configurable). Values are derived from OPNET default model.

ATM service category queue	Minimum guaranteed bandwidth (% of the link)
CBR	50%
RT-VBR	30%
NRT-VBR	10%
ABR	5%
UBR	5%

Table 8: Assignment of queue weight (configurable).

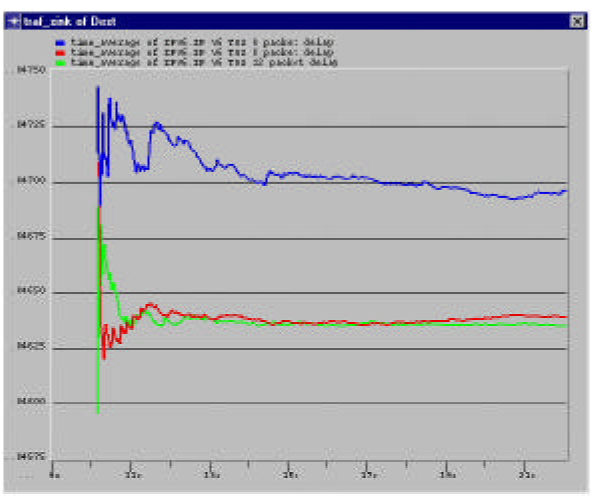

Figure 13: IPv6 packet delay (sec) vs. simulation time for five different ATM service categories assigned to three ToS values (0, 8, and 12).

5 Conclusion

Through the use of OPNET simulation tool, we demonstrated the upgrade of the IP protocol in which an IPv6-over-ATM network is capable of producing services guarantees that IP alone could not deliver. The simulation results indicated the feasibility of mapping high priority IPv6 packets (specified by the ToS value in the packet header) to a better QoS ATM SVC's when serving application that requires tighter end-to-end packet delays. By taking advantage of the ATM virtual circuit technology with unique connection identifiers (VPI/VCI) and a wide range of QoS offered by various ATM service categories, one could use the IPv6 ToS field in an IPv6-over-ATM network and offer a larger variety of QoS to different user applications.

Acknowledgment

The authors thank N. Alborz, B. Chen, and M. Nikolic for valuable comments and suggestions in preparation of the manuscript. This research was funded by the Canadian Foundation for Innovation Grant 910-99.

References

[1] "IP next generation overview," http://playground.sun.com/pub/ipng/html/INET-IPng-Paper.html.

[2] "IPv6: The new Internet protocol," http://www.comsoc.org/pubs/surveys/stallings/stallings-orig.html.

[3] S. Deering and R. Hinden, "Internet protocol, version 6 (IPv6) specification," RFC 2460, Dec. 1998 and RFC 1883, Dec. 1995: http://www.ietf.cnri.reston.va.us/rfc/rfc2460.txt, and http://www.ietf.cnri.reston.va.us/rfc/rfc1883.txt.

[4] "Asynchronous transfer mode (ATM) switching," http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/atm.htm.

[5] G. Armitage, M. Jork, P. Schulter, and G. Harter, "IPv6 over ATM networks," RFC 2492, Jan. 1999: http://www.ietf.cnri.reston.va.us/rfc/rfc2492.txt, or ftp://ftp.isi.edu/in-notes/rfc2492.txt.

[6] M. Hassan and M. Atiquzzaman, *Performance of TCP/IP Over ATM Networks*. Norwood, MA: Artech House Publishers, 2000.

[7] "OPNET ATM model description," http://www.opnet.com/products/library/ATM_Model_desc.pdf.