An Evaluation of RSVP Control Message Delivery Mechanisms

O. Komolafe and J. Sventek
Dept. of Computing Science, University of Glasgow, Glasgow G12 8QQ, UK.
femi@dcs.gla.ac.uk, joe@dcs.gla.ac.uk

Abstract— **This paper compares a number of RSVP control message delivery mechanisms in terms of several performance metrics. It is found that, unsurprisingly, the default best effort delivery of RSVP messages performs poorly in the face of increasing network congestion. This poor performance may be ameliorated by reserving bandwidth for RSVP messages or by implementing a recently proposed exponential back-off retransmission algorithm. The relative performance of the different control message delivery mechanisms is quantified and explained. For further insight, the impact of the induced loss of Path or Resv messages is also investigated. Interesting trends are observed and explained. A key discovery is that the loss of Path messages is typically more detrimental than the loss of Resv messages. This finding, which suggests that, relatively speaking, it is more important to deliver Path messages reliably than Resv messages, is attributed to the dependency of Resv states on Path states mandated by the RSVP standards.**

I. Introduction

Although originally proposed for resource reservation in Internet Protocol (IP) integrated services networks [1], the resource reservation protocol (RSVP) [2], [3] has been attracting considerable attention for use in multi-protocol label switching (MPLS) networks [4]. The motivation is to use RSVP as a signaling protocol to establish label-switched paths (LSPs) [5]. Another signaling protocol that may be used in MPLS networks is the constraint-based routing label distribution protocol (CR-LDP) [6], [7], which has been compared with RSVP [8]. Recently, MPLS signaling protocols, along with IP routing protocols, are being standardised as a control plane for packet, time-division, wavelength and spatial switching networks. Generalised MPLS (GMPLS) [9], [10], as this control plane is known, will therefore potentially use a suitably enhanced form of RSVP [11].

Hence, it is reasonable to conclude that RSVP will have a key role to play in future IP-centric networks, albeit in a different form and role than originally envisaged. Given the increase in the variety and importance of uses for RSVP, it is becoming increasingly important to ensure that RSVP control messages are reliably delivered. Herein lies a key distinction between RSVP and CR-LDP; the latter uses TCP transport to ensure that key control messages are reliably delivered whereas RSVP control messages are traditionally sent as best effort packets, with the expectation that periodic refresh messages will deal with occasional message loss. This paper evaluates a number of alternative control message delivery mechanisms for RSVP.

II. RSVP Overview and Motivation

The primary control messages in RSVP, Path and Resv messages, originate from the senders and receivers respectively. Path messages follow the route computed by the routing protocol and provide receivers with the description of the sender and traffic flow. Upon receipt of a valid Path message, each intermediate RSVP-capable router uses the information contained in the Path message to update or create a Path state entry for the sender before forwarding the appropriately updated Path message towards the receiver. After receiving a Path message, the receiver may make a reservation by sending a Resv message back to the source. The frequency with which Path and Resv messages are sent thereafter is determined by the soft-state refresh period. Although the Path state and the Resv state will eventually timeout if not refreshed, PathTear or ResvTear messages may be used to tear down state promptly.

The research community has focused on RSVP performance evaluation and improvement [12], [13], [14], and the standards have evolved accordingly. An early significant proposed enhancement to RSVP was to introduce bundle and summary refresh messages in an attempt to make RSVP more scalable [15]. Furthermore, in order to to improve the reliability of control message delivery, a message identification object, Acknowledgement messages and a retransmission algorithm were defined [15]. RSVP with traffic engineering extensions (RSVP-TE) [5] was later introduced to facilitate LSP establishment. Consequently, a label request object and a label object must be included in Path and Resv messages respectively. Recently, extensions to RSVP-TE have being defined for GMPLS signaling [11]. Naturally, the existing label-related objects must be augmented to allow packet, time-division, wavelength or spatial-switching connections to be established.

It is obviously desirable that RSVP control messages are reliably delivered. It was originally suggested that control messages be sent as best effort packets, with periodic refresh messages dealing with the occassional loss of control message or, alternatively, that some bandwidth be reserved for RSVP messages to protect them from congestion loss [3]. The different extensions to RSVP have broadly left the delivery of control messages unaltered, with the exception of RFC 2961 [15]. Based on earlier findings [12], RFC 2961 introduced a retransmission algorithm to improve the reliability of control message delivery. Essentially, when a node sends an initial, or trigger, Path or Resv message, it will indicate that it requires an acknowledgement by setting a flag in the message. The node retransmits the message, at an interval determined by an exponential back-off timer, while awaiting the corresponding Acknowledgement message or until a retry limit has been reached. Thereafter, refresh RSVP messages are sent as normal. Interestingly, RFC 2961 comments that the use of this exponential back-off procedure is worthy of further investigation. This paper seeks to evaluate the performance of the proposed retransmission algorithm relative to the more traditional methods of delivering RSVP control messages.

Four different control message delivery mechanisms are compared in this work:

- Sending RSVP messages as best effort packets.
- Reserving bandwidth for RSVP messages.
- Use of the exponential back-off retransmission algorithm for trigger Path and Resv messages.
- Use of the exponential back-off retransmission algorithm for all (i.e. trigger and refresh) Path and Resv messages.

These control message delivery mechanisms are compared in terms of a number of key performance metrics as the network congestion increases. Congestion will likely lead to the simultaneous loss of Path and Resv messages, making it difficult to identify the impact of losing either message type. Consequently, in order to gain further insight into the behaviour of these mechanisms, their performance in the face of induced losses of either Path or Resv messages is also investigated.

III. Approach

Network Simulator (NS) [16] is a widely used tool for studying data networks. An RSVP simulator, RSVP/ns [17], has been implemented in NS. For this work, RSVP/ns was enhanced to implement the control message delivery reliability-enhancing schemes proposed in RFC 2961 and was modified to lose Path and Resv packets randomly with a variable probability.

Figure 1 shows the exemplar topology considered. Initially, the sender sends a Path message to the receiver and simultaneously starts sending it a constant bit rate (CBR) data stream. Upon receipt of the Path message, the receiver generates a corresponding Resv message which is sent hop-by-hop back to the sender, with the intermediate routers reserving the appropriate resources for the flow. Thereafter, refresh Path and Resv messages are sent regularly to maintain the reservation. The links between the routers are saturated with exponentially distributed ON/OFF background traffic. Hence, when no reservation is made for the CBR flow it will contest for resources with the background traffic. The link bandwidths were set such that the aggregate traffic exceeds the link capacities, hence a proportion of the CBR and background traffic must be buffered and possibly lost. However, if a reservation is successfully made for the CBR flow, only the background traffic is buffered and potentially dropped. Therefore, monitoring appropriate performance metrics for the CBR flow gives insight into the ability of the control message delivery mechanism to establish and maintain reservations.

The key performance metrics considered are the throughput, the proportion of transmitted packets successfully delivered to the receiver, the mean end to end delay experienced by packets in the CBR flow and the protocol overhead, approximated using the total number of RSVP messages transmitted.

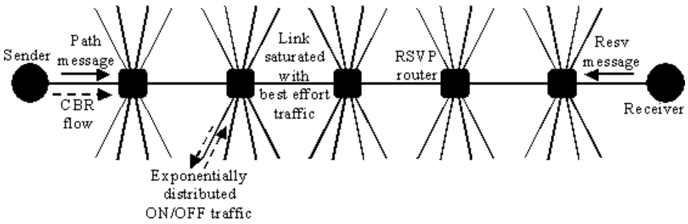

Fig. 1. Exemplar topology considered

RSVP timing values are important and the following recommended default values are used throughout this work:

- The lifetime expiration factor, the number of consecutive refresh messages which may be lost before a state timeout occurs, is 3 [3].
- The refresh interval for Path and Resv messages is 30s, although in order to avoid message synchronisation [18], the actual refresh period is randomly set between 15s and 45s, as recommended [3].
- The retransmission limit, the maximum number of times any message is retransmitted, is 3 [15].
- The initial retransmission interval is 2s, the round-trip interval between adjacent nodes [15].
- A power of 2 exponential back-off is used, hence the retransmission interval doubles with each attempt until the retransmission limit is reached [15].

Each exponentially distributed ON/OFF source generates 500byte packets at a rate of 500kb/s during ON periods. The mean ON and OFF periods were both 1s. The CBR source emitted 500byte packets at a rate of 500kb/s. The link delay was nominally set to 1s and the buffer size to 1000 packets. When investigating the impact of congestion on the performance of the different control message delivery mechanisms, the dimensions of the links were varied accordingly. On the other hand, the effect of the induced loss of Path and Resv messages was studied by varying the probability of losing each message type, with the link bandwidth set such that the background best effort traffic exactly saturates the link on average. 1kb/s was reserved for RSVP control messages when considering the impact of reserving bandwidth for RSVP messages.

IV. Impact of Congestion

A. Throughput

In order to make the simulation results statistically credible, a large number of simulation runs, each with different random number generation seeds, were carried out. The appropriate number of simulation runs was determined empirically. For example, when the relative link bit rate is 0.5 and the control messages are sent as best effort packets, it was found that the mean throughput for 10, 20, 30, 50 and 125 simulation runs are 0.489, 0.546, 0.578, 0.561 and

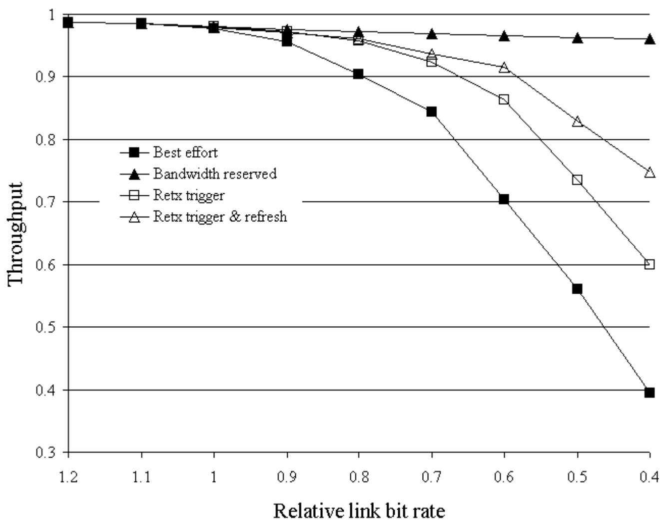

Fig. 2. Impact of network congestion on throughput

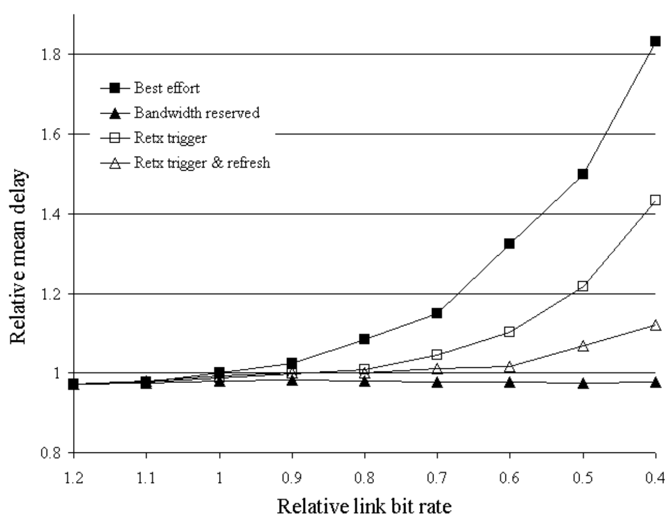

Fig. 3. Impact of network congestion on end to end delay

0.554 respectively. These findings, which suggest that 50 runs suffice, were found to be representative. Thus, the results presented in this paper are the averages of at least 50 simulation runs.

Throughput is affected by increasing congestion as depicted in Figure 2. Reserving bandwidth for the transmission of RSVP control messages always results in very high throughput. Sending the messages as best effort packets leads to the throughput falling with congestion, due to the increased delay and loss experienced by RSVP messages making it more difficult to establish and sustain reservations. Applying the retransmission algorithm to all unacknowledged control messages gives a higher throughput than applying the algorithm to trigger messages exclusively. This observation can be explained by realising that, under high congestion, sufficient successive refresh messages may be lost such that a Path or Resv state timeout occurs. Consequently, retransmitting refresh messages while awaiting the corresponding Acknowledgement messages results in the state timeout occurring less frequently, so a higher throughput is obtained.

B. Mean Delay

The impact of increasing congestion on the mean end to end delay experienced by CBR flow packets is shown in Figure 3. The largest delays are obtained when RSVP messages are sent as best effort packets. The lowest delays are obtained when bandwidth is reserved for RSVP messages. Retransmitting unacknowledged refresh and trigger messages produces a lower mean delay than when only unacknowledged trigger messages are retransmitted.

C. Protocol Overhead

The total number of RSVP messages sent is used to approximate the protocol overhead as it is indicative of the number of control messages that must be processed by RSVP nodes and the total network bandwidth consumed by RSVP. Figure 4 presents the number of RSVP messages transmitted for different levels of network congestion. It can be seen that, unsurprisingly, retransmitting both unacknowledged trigger and refresh Path and Resv messages produces significantly more RSVP messages than any other control message delivery mechanism and, interestingly, as the congestion increases the number of RSVP messages sent shows a slight rise before eventually falling. This rise is due to the increased number of retransmissions in an attempt to combat the growing number of control messages lost and delayed due to the increasing congestion. The number of messages remains largely unchanged when only trigger messages are retransmittable and when bandwidth is reserved for the control messages. In contrast, the number of RSVP messages drops when the control messages are sent as best effort packets. Naturally, the source will still send the same number of Path messages, regardless of the network congestion. However, the aggregate number of RSVP messages exchanged is dependent upon the number of nodes that have Path or Resv states established in them, a number evidently dependent on the network congestion.

V. Impact of Induced Message Loss

The performance of the different control message delivery mechanisms in the face of increasing network congestion has been evaluated in the preceding section. Using the same performance metrics, the impact of the induced loss of either Path or Resv messages on the ability of RSVP to establish and sustain reservations was investigated.

A. Throughput

Figure 5 shows the impact of an increasing Path or Resv message loss probability on the throughput for each control message delivery mechanism. It is evident that, in general, sending control messages as best effort packets leads to the lowest relative throughput, followed by reserving bandwidth for the delivery of control messages. Retransmitting only the unacknowledged trigger messages yields

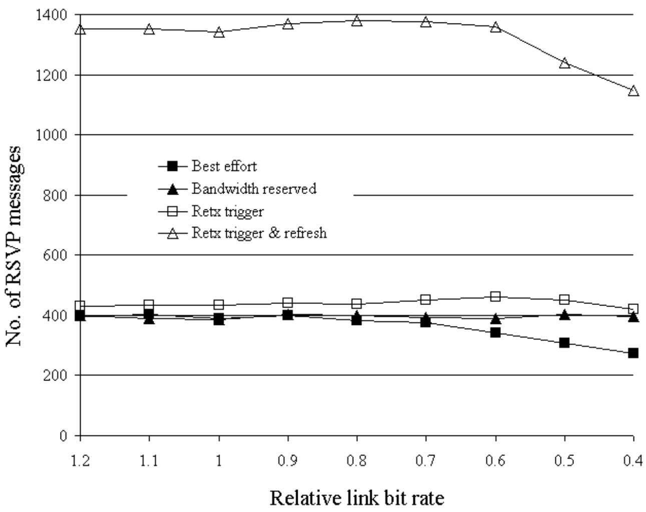

Fig. 4. Impact of network congestion on protocol overhead

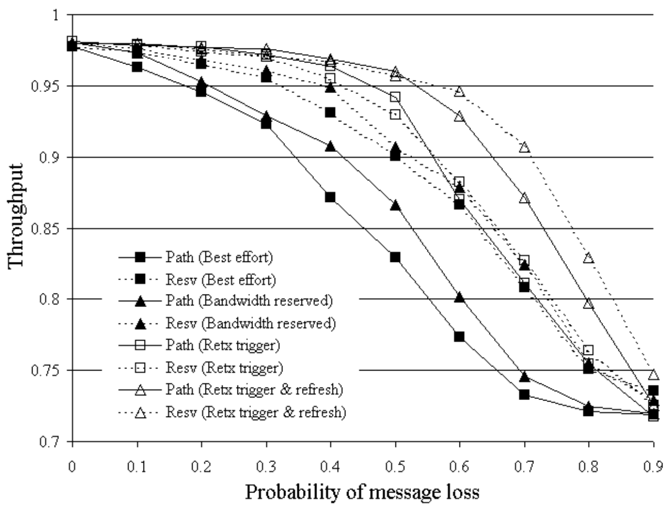

Fig. 5. Impact of message loss on throughput

the second highest throughput. The highest throughput is obtained when the retransmission of refresh and trigger messages is allowed. For high message loss probabilities, the throughput approaches the value obtained when the CBR flow is sent as best effort traffic, 0.7175. Hence, the obtained throughput values may be compared with 0.7175 to indicate the ability of the control message delivery mechanism to establish and sustain reservations. It can be seen from Figure 5 that, regardless of the control message delivery mechanism, the loss of Path messages is more detrimental than the corresponding loss of Resv messages. It also appears that the results for the loss of Path messages are more dependent on the control message delivery mechanism than the loss of Resv messages since significantly different results are obtained for each series when Path messages are lost. In contrast, only the retransmission of trigger and refresh messages produces noticeably different results when Resv messages are lost; the other three control message delivery mechanisms produced comparable results, especially when the message loss probability is relatively high.

B. Mean Delay

The effect of the loss of Path and Resv messages on the mean end to end delay experienced by packets in the CBR flow is illustrated in Figure 6. As previously observed, the loss of Path messages is more detrimental than the corresponding loss of Resv messages and the meritocratic ordering of the control message delivery mechanisms is retransmitting unacknowledged trigger and refresh messages, retransmitting only unacknowledged trigger messages, reserving bandwidth for the delivery of RSVP messages and, lastly, sending the control messages as best effort packets. As the message loss probability increases, the mean delay obtained converges to the relative mean delay for the CBR flow when no reservation is made, 1.846. Figure 6 also suggests that, as observed from Figure 5, the control message delivery mechanism is more consequential when Path messages are lost than when Resv messages are lost.

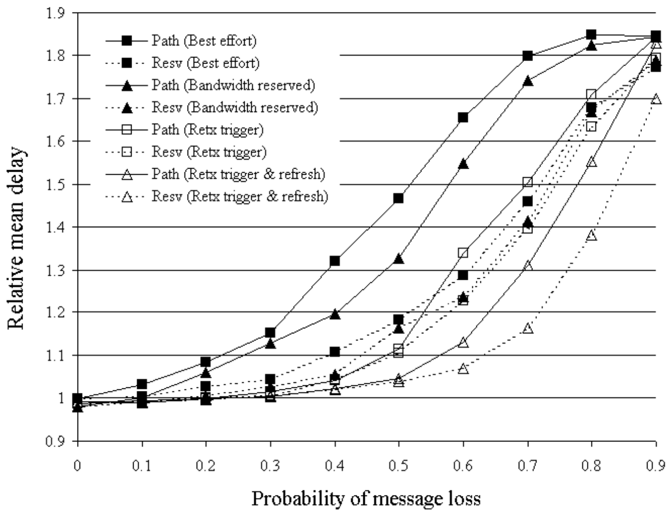

Fig. 6. Impact of message loss on end to end delay

C. Protocol Overhead

Figure 7 shows that, as noted from Figure 4 earlier, retransmitting unacknowledged trigger and refresh messages results in the greatest number of RSVP messages being transmitted, followed by retransmitting unacknowledged trigger messages. Sending the messages as best effort packets or reserving bandwidth exclusively for RSVP messages result in a similar number of messages being sent

However, perhaps the most striking feature of Figure 7 is that the loss of Path and Resv messages impact the total number of RSVP messages sent in a significantly different manner. When Resv messages are lost, the total number of RSVP messages sent drops in a linear fashion, as would be expected. However, the loss of Path messages results in a sub-linear decrease in the number of RSVP messages sent. This finding suggests that the transmission of other RSVP control messages is dependent upon Path messages, a deduction that may be easily verified by realising that, for

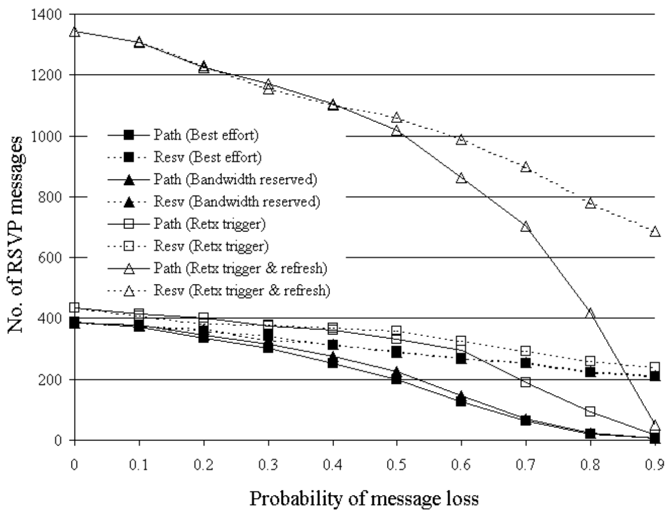

Fig. 7. Impact of message loss on RSVP overhead

example, the receiver cannot send the first Resv message until it receives the first Path message.

For all the performance metrics considered thus far, it has been observed that the loss of Path messages is more detrimental than the loss of Resv messages and that the control delivery mechanism affects the performance metric more when Path messages are lost than when Resv messages are lost. The rules for RSVP require that when a Path state timeout occurs, the Path state *and* any dependent Resv state(s) be deleted [3], [19]. In contrast, a Resv state timeout only results in the Resv state being deleted. In order words, Resv states are dependent on Path states yet Path states are not dependent on Resv states. This fact suggests maintaining Path states at nodes is, relatively speaking, more important than maintaining Resv states.

VI. Conclusions & Future Work

RSVP has evolved from being a protocol for reserving resources in the Internet to being considered as a generic signaling protocol for the GMPLS control plane. As such, the reliable delivery of RSVP control messages is of ever-increasing importance. The initial default of sending the control messages as best effort packets has been found wanting with suggested solutions including reserving bandwidth exclusively for RSVP control messages or the recently proposed exponential back-off retransmission algorithm. This paper has compared a number of alternative control message delivery mechanisms in terms of their ability to cope with increasing network congestion. It was found that, unsurprisingly, reserving bandwidth for RSVP messages obtains the best performance and sending the messages as best effort packets the worst. The proposed retransmission algorithm performs between these two extremes; its performance may be improved by retransmitting all unacknowledged messages, rather than only unacknowledged trigger messages. The performance of the alternative control message delivery mechanisms in the face of the induced loss of Path or Resv messages was also investigated. It was discovered that the loss of Path messages is more detrimental than the loss of Resv messages. This finding was attributed to the dependency of Resv states at nodes on the corresponding Path states mandated by the RSVP standards. This finding suggests that, relatively speaking, more attention should be paid to the reliable delivery of Path messages than Resv messages.

It will be interesting, therefore, to study the performance of such asymmetric control message delivery mechanisms as they may perform similarly to the symmetric schemes evaluated, at much lower protocol overhead. Another interesting area of future work is to determine the findings' sensitivity to RSVP timing values. Such studies will provide guidance regarding expedient timing values for RSVP. Furthermore, maximising RSVP's performance while minimising, say, the protocol overhead is an intriguing multi-objective optimisation problem.

References

[1] R. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, June 1994.

[2] L. Zhang, S. Deering, D. Estrin, S. Shenker, D. Zappala, "RSVP: A New Resource Reservation Protocol", IEEE Network, Vol. 7, pp. 8-18, Sept. 1993.

[3] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin, "Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification", RFC 2205, Sept. 1997.

[4] E. Rosen, A. Viswanathan, R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, Jan. 2001

[5] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels', RFC 3209, Dec. 2001.

[6] L. Andersson, P. Doolan, N. Feldman, A. Fredette, B. Thomas, "LDP Specification", RFC 3036, Jan. 2001.

[7] B. Jamoussi, L. Andersson, R. Callon, R. Dantu, L. Wu, P. Doolan, N. Feldman, A. Fredette, M. Girish, E. Gray, J. Heinanen, T. Kilty, A. Malis, "Constraint-Based LSP Setup using LDP", RFC 3212, Jan. 2002.

[8] A. Ghanwani, B. Jamoussi, D. Fedyk, P. Ashwood-Smith, L. Li, N. Feldman, "Traffic Engineering Standards in IP Networks Using MPLS", IEEE Communications Magazine, pp. 49-53, Dec. 1999.

[9] A. Banerjee, J. Drake, J. Lang, B. Turner, K. Kompella, Y. Rekhter, "Generalized Multiprotocol Label Switching: An Overview of Routing and Management Enhancements", IEEE Communications Magazine, pp. 144-150, Jan. 2001.

[10] L. Berger (Ed.), "Generalised Multi-Protocol Label Switching (GMPLS) Signaling Functional Description" RFC 3471, Jan. 2003.

[11] L. Berger (Ed.), "Generalised Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions" RFC 3473, Jan. 2003.

[12] P. Pan, H. Schulzrinne, "Staged Refresh Timers for RSVP", Proc. IEEE Globecom, 1997.

[13] A. Neogi, T. Chiueh, P. Stirpe, "Performance Analysis of an RSVP-Capable Router", IEEE Network, pp.56-63, Sept./Oct. 1999.

[14] L. Wang, A. Terzis, L. Zhang, "A New Proposal for RSVP Refreshes", Proc. IEEE Int. Conference on Network Protocols, 1999.

[15] L. Berger, D. Gan, G. Swallow, P. Pan, F. Tommasi, S. Molendini, "RSVP Refresh Overhead Reduction Extensions" RFC 2961, April 2001.

[16] http://www.isi.edu/nsnam/ns

[17] http://www.isi.edu/nsnam/ns/ns-contributed.html

[18] S. Floyd, V. Jacobson, "Synchronisation of Periodic Routing Messages", IEEE/ACM Transactions on Networking, Vol.2, No.2, pp. 122-136, April 1994.

[19] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin, "Resource ReSerVation Protocol (RSVP) - Version 1 Message Processing Rules", RFC 2209, Sept. 1997.