

QUANTITATIVE PREDICTION OF NACK-ORIENTED RELIABLE MULTICAST (NORM) FEEDBACK

R. Brian Adamson
Joseph P. Macker
Information Technology Division
Naval Research Laboratory

ABSTRACT

We have applied the concept of truncated exponential timers for efficient reliable multicast feedback suppression for cases of both multicast and unicast feedback channels. Unicast feedback operation for multicast transport is becoming a more prevalent concern with the advent of source specific multicast routing and asymmetric networks offering forward-based multicast (e.g., satellite distribution network). We discuss our approach to the design and its integration with a working reliable multicast protocol. We then present simulation results demonstrating that observed implementation performance matches the analytically predicted performance. Finally, we formulate a quantitative predictor of reliable multicast protocol feedback traffic levels.

BACKGROUND

Applications requiring group communication can benefit from the basic Internet Protocol (IP) suite multicast model [Deering89] that provides network layer multipoint delivery of group-addressed packets. The IP multicast model uses generic IP datagrams as its raw service and does not provide inherent reliability of data delivery. IP multicast applications requiring guaranteed delivery need reliable multicast (RM) transport mechanisms [Mankin98]. Reliable bulk-transfer, multicast transport is a desirable technology for distribution of data to a group on the Internet.

Previous work has shown that selective Negative-Acknowledgement (NACK) reliable multicast protocols are more scalable than those based on positive acknowledgement (ACK) of received data [Pingali93]. NACK-oriented reliable multicast (NORM) protocols offer scalability advantages for applications and/or network topologies where, for various reasons, it is prohibitive to construct a higher order delivery infrastructure above the basic Layer 3 IP multicast service (e.g. unicast or hybrid unicast/multicast data distribution trees). Additionally, the scalability property of NACK-oriented protocols [Levine96] may be applicable where broad "topological fanout" is expected for a single network hop (e.g. cable-TV data delivery, satellite, or other broadcast communication services). Furthermore, the simplicity of a protocol based on "flat" group-wide multicast distribution may offer advantages for a broad range of distributed services or dynamic networks and applications.

To avoid NACK feedback implosion (feedback traffic levels well in excess of the network's or sender's capacity), protocols have been designed to suppress redundant NACK transmissions among a group of receivers using probabilistic techniques (e.g., SRM [Floyd95], MDP [Macker99]). In these protocols, receivers use random back-off timers to delay repair request transmission, sending their NACK only if it is not superseded by an overriding repair request from another receiver. In past work, these strategies have generally assumed reciprocal multicast connectivity among the multicast session participants and have used uniform probabilistic backoff window techniques.

Deployment of IP multicast services can be complex and has been slow in its proliferation to the Internet outside of research networks. Network management and operation for multicast requires additional consideration in terms of device configuration complexity, address management, and security. Also, in some cases, where the network topology provides asymmetric connectivity with broadcast connectivity predominantly available in only one direction, deployment of non-reciprocal multicast operation is more feasible than providing classical "many-to-many" multicast routing services. The development of single source multicast (SSM) protocols [Holbrook99] within the Internet Engineering Task Force (IETF) is being done in anticipation that SSM will simplify and speed the deployment of multicast services to networks, particularly for bulk transfer applications. The challenge for NORM is to remain scalable in such environments where multicast transmission may be available from a sender to a group, but the receivers are restricted to provide feedback via unicast connectivity paths.

NORM requires feedback from receivers in the form of requests (i.e. NACKs) indicating the receivers' needs for retransmissions or repairs to complete reliable transport. The amount of feedback would grow linearly with the group size for protocols that do not address this issue. Approaches exist to avoid feedback implosion including timers, tokens, and hierarchies. Hierarchical approaches do not fit within the goals of a "flat" topology NORM protocol and the use of tokens can be of limited scalability or lead to large latency in feedback. In addition, both hierarchies and token-based schemes may be difficult to apply to more dynamic environments such as those in wireless-based networks. An approach has been developed to use timer-based feedback suppression to facilitate scaling of NORM protocols for both networks with ubiquitous multicast connectivity and those limited to unicast receiver feedback.

APPROACH

Exponential Timer Based Feedback Suppression

The use of timer-based feedback suppression can be accomplished in an end-to-end fashion and is thus adaptable to a wide variety of networks with few assumptions on topology. The relative performance of feedback suppression backoff timers based on an exponential distribution is described in [Nonnemacher98]. This approach can readily adapt to the group size, maintaining low-latency feedback even with only a rough estimate of actual group size [Fuhrmann01],[Widmer01]. While it is expected that this approach can minimize the amount of expected feedback for timer-based feedback suppression, deployment of a real protocol requires some quantitative prediction of its impact on the network.

As a framework for this study, we modified the Multicast Dissemination Protocol (MDP), a NACK-oriented RM protocol, to use exponential suppression timers. While we performed work studying performance within a specific, existing, operational RM protocol, we feel our findings are general in nature. The approach to determine the applicability and performance of MDP NACK suppression was to develop an analytical model to predict suppression performance and compare it to results obtained via network simulation as the multicast group size was increased. The actual MDP implementation was designed to operate in simulation environments (i.e., ns2 and OPNET) as well as in real systems. Most of the studies described here were performed within the ns2 simulation environment [ns].

The suppression timers we explored were randomly scheduled using a truncated exponential distribution with the density function:

$$f(x) = \begin{cases} \frac{1}{e^\lambda - 1} \cdot \frac{\lambda}{\tau} e^{-\frac{\lambda}{\tau} x} & , 0 \leq x \leq \tau \\ 0 & , \text{otherwise} \end{cases} \quad (1)$$

MDP receivers use this density function to schedule random backoff timeouts over the finite interval $[0, \tau]$. In MDP, the scaling factor λ is set according to the optimization given in [Nonnemacher98] which is:

$$\lambda_o = \ln(R) + 1 \quad (2)$$

where R is the estimated number of receivers providing feedback. Given R , this technique for scheduling backoff timeouts is implemented in MDP receivers using the following algorithm:

- 1) Pick a uniformly distributed random number x in the range of:

$$\frac{\lambda}{\tau(e^\lambda - 1)} \text{ to } \frac{\lambda}{\tau(e^\lambda - 1)} + \frac{\lambda}{\tau}$$

- 2) This uniform random variate can be transformed to an exponential one [Abramowitz70] to generate the desired random backoff time with the following equation:

$$t_{backoff} = \frac{\tau}{\lambda} \ln \left(x \cdot (e^\lambda - 1) \cdot \frac{\tau}{\lambda} \right) \quad (3)$$

In MDP, the backoff window interval τ is picked as a multiple of the sender's current Group Round Trip Time (GRTT) estimate collected within the protocol such that $\tau = n \cdot GRTT$. In the current MDP implementation, n can be set to greater than or equal to one. The sender collects round trip measurements from timestamp information on feedback packets from the receiver set. The MDP GRTT value is a filtered estimate of the maximum round trip time between the sender and any receiver in the group. Within MDP, the GRTT is advertised to the receiver set so that they may properly scale feedback backoff and other protocol timeouts. Setting τ to a desired bound based on GRTT can control the maximum delay experienced in triggering feedback message events.

Suppression Performance for Multicast Feedback Messages

According to [Nonnemacher98, eq 16], τ can be set as a multiple of the one-way receiver-to-receiver delay to produce a desired expected number $E(x) = N$ of transmitted receiver feedback messages for a given feedback event using

$$k = 1.2 \frac{\lambda_o}{\ln(N)} \quad (4)$$

where λ_o is the optimal λ from Equ. (2) for a given group size R , and k is the multiplier of the one-way receiver-to-receiver delay. Note k must be greater than one to begin yielding any suppression. Otherwise, all receiver backoff timers will expire before any suppressing messages can arrive from other receivers.

A metric quantifying the effectiveness of feedback suppression which we term the *suppression factor* (SF) can be defined as

$$SF = N/R \quad (5)$$

which is the ratio of the expected number of feedback messages N to the total receiver set size R . This metric is useful because it can also be empirically measured in simulated and practical protocol applications. An analytic prediction of the *suppression factor* for multicast feedback suppression based on exponentially distributed backoff timeouts can also be formulated. Substituting the formula given by Equ. (2) for λ_o , Equ. (4) can be transformed to predict the expected number of multicast feedback messages for an event given a group size R :

$$N = e^{\frac{1.2}{k}(\ln(R)+1)} \quad (6)$$

For a topology with homogeneous delays, the MDP GRTT time is two times the one-way delay and so k can be replaced with $2 \cdot n$ where n is the multiple of GRTT comprising the maximum backoff window. Thus, given that the suppression factor is defined as the ratio of N/R , an analytical

approximation of SF_m for multicast feedback can be given by:

$$SF_m = \frac{e^{\frac{1.2}{2n}(\ln(R)+1)}}{R}; n > 1 \quad (7)$$

This model of expected feedback suppression performance is based on an assumption of equal delays among the receivers. However, note that such homogeneous delay is a worst-case scenario for feedback suppression performance, so this model establishes some measure of expected worst-case performance. As the delays become heterogeneous, the suppression performance actually improves slightly as suppressing feedback messages are delivered earlier to some receivers, thus increasing the probability of suppression. This is reinforced in MDP by use of the maximum round trip delay (GRTT) as the reference time period for scaling the feedback backoff interval.

Suppression Performance for Unicast-to-Source Feedback Messages

Multicast feedback from the receiver set may not always be possible in some networks. Unicast-to-source feedback is potentially likely in large fan-out satellite networks or within future SSM multicast sessions so we consider it an important area for analysis and consideration. Suppression of unicast

feedback messaging can be implemented by having the sender retransmit (via multicast) received feedback messages to the group as they are received or alternatively advertise the cumulative state resulting from the received feedback. MDP senders suppress unicast NACKing using repeated, but controlled, advertisement of repair state in response to received NACKs.

The logical network topology illustrated in Fig. 1a represents a network with homogeneous delays between the sender and receivers. The assumption is that the links represented in the figure are bi-directional and have equal round trip delays. Note that this very closely corresponds to the real case of satellite or other one-hop broadcast networks. The analytical feedback suppression performance predictor given in Equ. (7) can be modified to predict performance for suppression of unicast-to-source feedback, again assuming homogeneous delays among the group members. To do this, it is simply necessary to realize that the receiver-to-receiver delay in delivery of feedback messages (or accumulated feedback state) is the driving factor to suppression performance along with the backoff timer distribution and window size. Figure 1 provides reference logical topologies for visualizing the receiver-to-receiver messaging delay for multicast and unicast feedback in different scenarios.

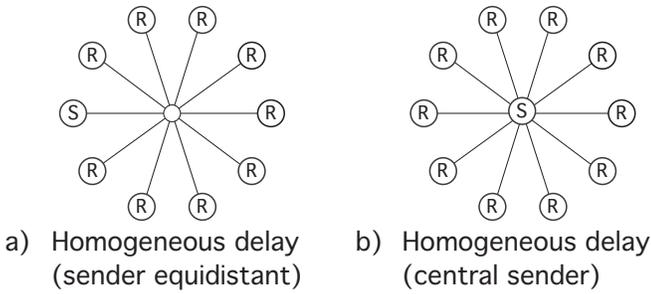


Fig. 1 – Example Logical Group Topologies

For multicast feedback suppression, the message delivery delay among receivers for this logical topology is equal to traversing two of the "links" illustrated in Fig. 1a which is equal to one-half of the round trip time (and GRTT in the case of MDP) among the group members. For unicast suppression, NACKs are logically forwarded through the sender node. Thus, the delay of the delivery of feedback information from a given receiver to other receivers becomes four link traversals (again referring to the logical topology in Fig. 1a). This doubles the logical "one-way" receiver-to-receiver delivery delay over the case of direct multicast connectivity among the receivers. So, to predict unicast feedback performance, Equ (7) can be modified by removing the factor of 2 applied to n to yield a predicted unicast suppression factor (SF_u):

$$SF_u = \frac{e^{\frac{1.2(\ln(R)+1)}{n}}}{R} \quad (8)$$

As with Equ. (7), this is based upon homogeneous delays with the sender on an equal basis with the receiver set resulting in an average receiver-to-receiver delivery delay (through the sender) no greater than one GRTT, since the sender/receiver round trip also involves 4 logical link traversals (See Fig 1a). Given that

the suppression timer window is based on GRTT, this logical topology also represents a worst-case situation for unicast feedback. It is interesting to note as the sender becomes more centrally located in the topology, the average receiver-to-receiver delay (through the sender) can reduce with respect to the GRTT controlled backoff window, thus increasing suppression performance for unicast feedback. In the ideal case with a centrally located sender, the performance of unicast feedback suppression can actually equal that of multicast as the receiver-to-receiver delay is governed by only 2 logical link traversals. An example of this type of homogeneous delay topology is illustrated in Fig. 1b.

Fig. 2 is a graph plotting the predicted feedback suppression factor versus the receiver group size. The suppression factors, S_m and S_u , from Equ. (7) and (8), respectively, were plotted versus group size R assuming a suppression backoff window of 4 times the GRTT ($n = 4$) to yield the plots for exponentially distributed feedback. For comparison, the performance of uniformly distributed multicast and unicast feedback cases were also plotted using [Nonnemacher98, eq 10], normalized by R .

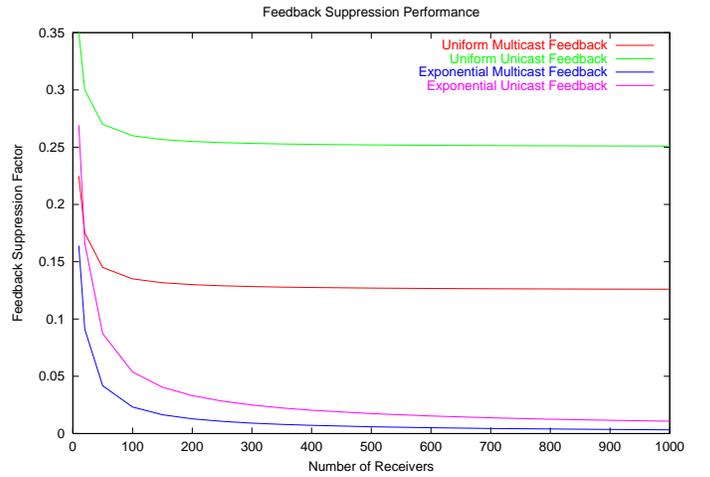


Fig. 2 – Predicted Feedback Suppression Performance

The *suppression factor* can be interpreted as the percentage of receivers actually providing feedback from all receivers needing to respond to an event (e.g. packet loss). From Fig. 2, it is easy to see how exponentially distributed timers result in a diminishing fraction of receivers providing feedback as the group size increases. Thus, the volume of non-suppressed feedback traffic grows very slowly in this case. With uniformly distributed timers, the feedback traffic volume grows linearly with the group size.

While the analytical results presented in Fig. 2 illustrate the case for use of exponential backoff timers and provide relative performance results for multicast and unicast feedback scenarios, this information does not provide a *quantitative* indication of feedback performance achieved for an actual protocol. However, the analytical model developed here is useful when combined with empirical results of actual protocol performance obtained in simulation or real world testing.

SIMULATION MODEL AND RESULTS

The goal of the work described in this paper is to provide a quantitative prediction of the scaling impact of applying exponential suppression techniques in a NORM protocol. As we have mentioned, we adopted MDP as an evaluation and implementation framework. The MDP protocol can be built in the *ns-2* discrete event network simulation tool [ns] and simulations can be instrumented to measure network loading as the protocol is run in realistic scenarios. Additionally, MDP receivers running in the simulation keep track of the number of NACK feedback messages sent and the number of potential NACKs that are suppressed from transmission. In this case the measured *suppression factor* (the ratio of transmitted feedback messages to total possible feedback messages) for the simulation can be directly calculated from the collected data as:

$$SF_{measured} = \frac{NACK_{Sent}}{NACK_{Sent} + NACK_{Suppressed}} \quad (9)$$

An *ns-2* simulation scenario was constructed to simulate the topology illustrated in Fig. 1a. The MDP protocol agents were configured to use a feedback backoff window of 4 times GR TT. Transmissions from the MDP sender node were subject to 10% random packet loss. This resulted in correlated loss among the receivers. Multiple simulation runs were conducted with MDP configured for combinations of unicast and multicast feedback with trials for uniformly distributed timer backoffs in addition to trials with exponentially distributed timer backoffs. The receiver set size was varied from as little as 10 nodes to 1000 nodes and the *suppression factor* was determined from the log of NACKs transmitted and suppressed by each receiver. The observed suppression factors for these trials are plotted along with analytically predicted suppression factors (SF_m and SF_u) in Fig. 3.

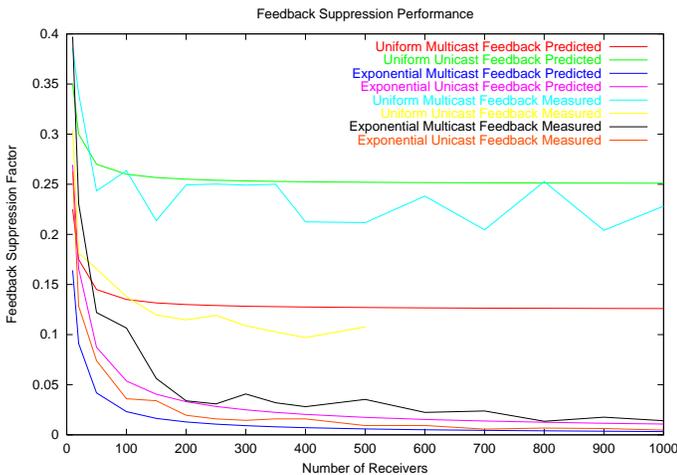


Fig. 3 – Measured Feedback Suppression Performance (10% correlated loss, data loss only, no FBM loss)

Note the simulations conducted to obtain the results shown in Fig. 3 were carefully constructed to evaluate MDP feedback performance under conditions matching the assumptions for the development of the analytical model as closely as possible. For example, while there was 10% correlated loss of data messages from the sender to excite NACK feedback messages (FBMs) from the receivers, the simulation was configured for no loss of

the generated FBMs. For the different trials conducted, MDP’s feedback suppression performance closely matched the predicted level of performance. Since the full protocol behavior is more complex than the assumptions underlying the analytical model, some differences were expected. More importantly, Fig. 3 shows that the *observed trend* of MDP’s suppression performance within the additional complexity of an implementation matches what is expected analytically, thus allowing MDP’s feedback demands to be favorably predicted for group sizes beyond the scale of simulation or test.

Fig. 4 presents the results for a similar set of trials with *uncorrelated* loss of the sender data among the receivers and again no loss of FBMs. The impact of uncorrelated loss appears as a linear shift in the measured values of the suppression factor metric, but once again the trend as the group size is increased directly corresponds to that of the analytical prediction. The use of packet-based forward error correction (FEC) for repair in MDP helps uncorrelated packet loss be treated the same as correlated loss with respect to the repair process. However differences in the exact quantity of packet loss for a given coding block among the receivers results in increased NACK messaging as some portion of early NACKing receivers have little loss and send non-suppressing NACKs.

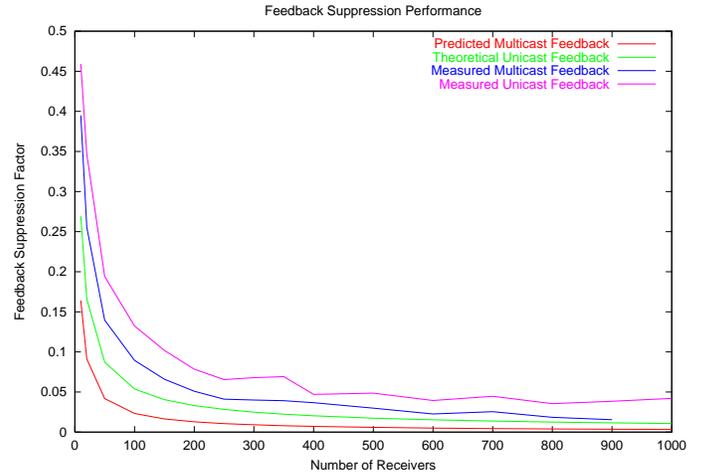


Fig. 4 – Measured Feedback Suppression Performance (10% uncorrelated loss, data loss only, no FBM loss)

Figure 5 illustrates the result of trials with correlated FBM loss as well as correlated sender data loss. The results indicate that the impact of FBM loss on multicast suppression performance is minimal; particularly as the group size grows large. The detriment to unicast feedback suppression is more pronounced as the process of relaying FBM information through the sender results in a single point of loss failure as any FBM loss event impacts the entire receiver group. The MDP protocol allows packing cumulative feedback information into the messages it multicasts to the group in response to received FBMs. This strategy helps diminish the impact of FBM loss under unicast feedback suppression. Again the decrease in the percentage of receivers providing feedback, matches the rate of decrease predicted by the analytical model as the group size is increased.

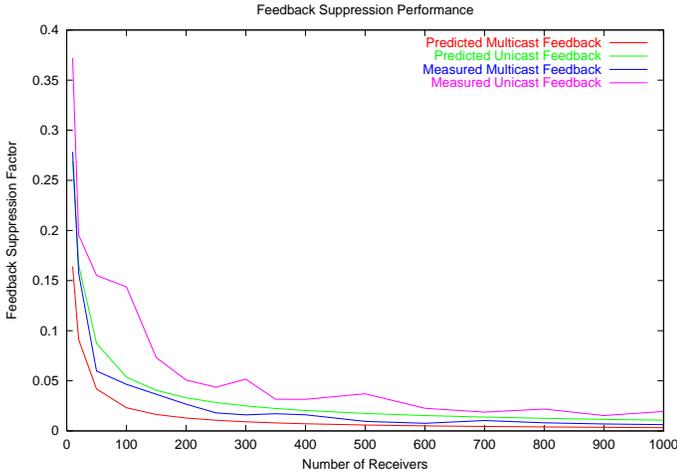


Fig. 5 – Measured Feedback Suppression Performance (10% correlated loss, all message types)

Finally, Fig. 6 depicts results collected from trials configured for uncorrelated loss of all message types including sender data and FBM traffic. The performance of multicast feedback is not significantly changed from that of Fig. 4. The performance of unicast feedback, as affected by the combination of uncorrelated loss and FBM loss, is only slightly diminished over the case of uncorrelated loss with no FBM loss or correlated loss with FBM loss. In all cases, the trends in suppression performance as the group size increases is consistent with that predicted by the analytical models.

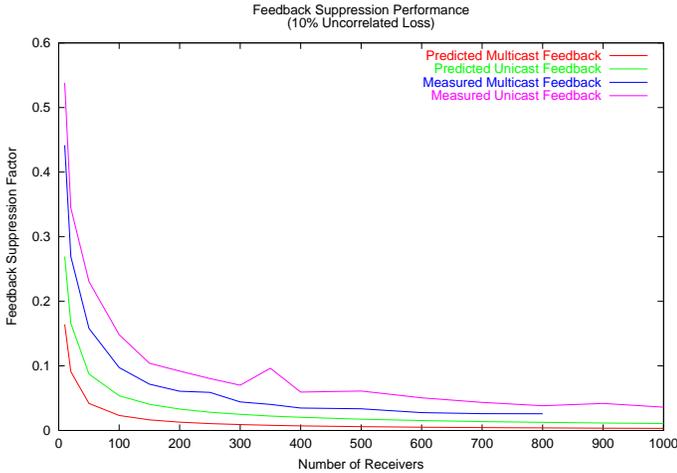


Fig. 6 – Measured Feedback Suppression Performance (10% uncorrelated loss, all message types)

Quantitative Feedback Traffic Level Prediction

Since the trend of the suppression performance of the implemented protocol follows that of the analytical predictor, we feel it is possible to estimate MDP feedback traffic loading for group sizes beyond the current capabilities of simulation or empirical evaluation. The quantity of feedback traffic for an anticipated group size can be predicted by first measuring the level of feedback ($f_{measured}$) for a given group size ($R_{measured}$). Then, note by Equ. (5), the feedback traffic level is linearly proportional to the group size and the feedback suppression factor for that group size:

$$f_{measured} \propto R_{measured} \cdot SF(R_{measured}) \quad (10)$$

For multicast feedback, Equ. (7) can be substituted for the suppression factor formula $SF(R_{measured})$ yielding:

$$f_{measured} \propto e^{\frac{1.2}{2n}(\ln(R_{measured})+1)} \quad (11)$$

Since the expected feedback traffic volume for a anticipated group size is similarly proportional to the corresponding suppression factor, the ratio of the anticipated feedback traffic level to the measured feedback traffic level can be expressed as:

$$\frac{f_{anticipated}}{f_{measured}} = \frac{e^{\frac{1.2}{2n}(\ln(R_{anticipated})+1)}}{e^{\frac{1.2}{2n}(\ln(R_{measured})+1)}} \quad (12)$$

Then, this expression can be further reduced to:

$$f_{anticipated} = f_{measured} \cdot \left(\frac{R_{anticipated}}{R_{measured}} \right)^{\frac{1.2}{2n}} \quad (13)$$

This provides a quantitative prediction of the feedback traffic level for the anticipated group size. Note that as the group size increases, the increase in the level of feedback traffic is relatively small. The size of the backoff window can be set conservatively (larger value of n) to limit the rate of feedback traffic level increase as the group size grows. This is an important observation because we wish the protocol to maintain high suppression factors even when the group estimate is inaccurate. To illustrate this, the level of cumulative feedback traffic for an MDP simulation with 500 receivers was measured. Fig. 7 is a plot of feedback traffic rate versus time for the simulation.

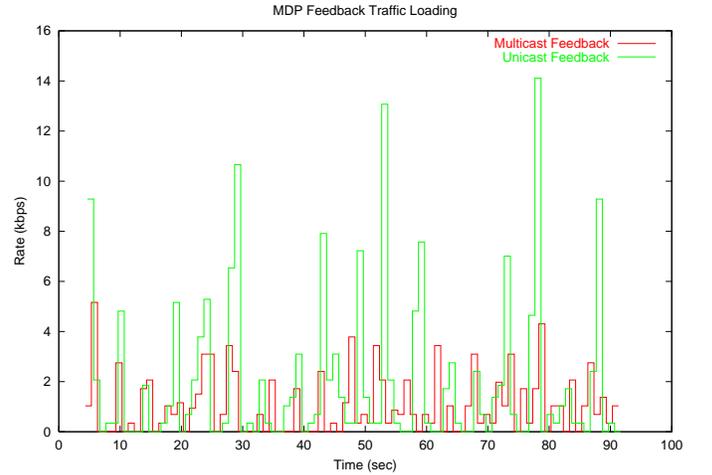


Fig. 7 – Measured MDP Cumulative Feedback - 500 Receivers (10% uncorrelated loss, all message types)

For multicast feedback messaging, the average cumulative feedback message loading was approximately two kilobits per second (kbps). Recalling that a backoff window equal to 4 times the GRTT ($n=4$) and using Equ. (13), it is possible to estimate that the average feedback message loading when the group size is increased to 1000 receivers would be approximately 2.2 kbps. For a group size of 10,000 receivers the predicted level of cumulative feedback traffic would be an approximate average of 3.1 kbps. In MDP, the quantity of NACK events is a principally a function of the sender transmission rate and packet loss rate. Since MDP receivers

initiate NACK repair cycles only on FEC code block boundaries, the sender FEC block size also impacts the quantity of NACK events. For the simulations conducted, the sender rate was 32 kbps with a packet segment size of 512 bytes and 30 packets per FEC coding block. The relatively low data rate was selected to enable the simulation to scale to the large group sizes evaluated. Given the impact of the FEC coding block size and the average packet loss rate, MDP feedback traffic levels as the sender's transmission rate is changed could also be approximated. Further work could be conducted to evaluate the specific quantitative impact of changes in these protocol parameters.

CONCLUSIONS

We have applied and studied the concept of truncated exponential timers for efficient timer-based RM feedback suppression. We developed a set of analytical predictors of feedback traffic levels for an actual working peer-to-peer RM protocol implementation. We also introduced simulation models and collected results under various loss and basic topology scenarios to compare against our analytical expectations. The simulation results we collected and examined show good agreement with analytical expectations. We believe this allows for reasonable confidence in using the analytical models as predictors in more scaled scenarios.

We presented results, developed models, and produced working implementations for cases of both multicast and unicast feedback channels. With the advent of SSM and satellite multicast architectures we feel the unicast feedback channel suppression improvements are important and that our results indicate that a high degree of suppression can be achieved. The MDP protocol with the mechanisms described in this paper has also been successfully deployed over a large-scale operational multicast satellite network in which the feedback channel is an important resource to conserve.

We feel our results help state a strong case for the use of truncated exponential timers within NORM protocols and we have demonstrated that the predicted trends hold up within an actual complete protocol implementation. It is anticipated that there will be increased deployment of multicast data distribution using MDP or similar protocols. For mission critical military systems, it is important that such deployments do not jeopardize network operation and can be expected to perform well in their anticipated deployment. The methodologies developed in this paper allow for quantitative prediction of the impact to the network of NACK-oriented reliable multicast deployment. Additionally, the techniques developed here will be used in conjunction with the ongoing development of multicast congestion control techniques [Adamson01], [Handley01] which will allow NORM protocols to be compatibly deployed as part of existing IP networks. The results and methodologies given here will be applicable to the evaluation of the performance of such reliable multicast protocol mechanisms in addition to NACK which require feedback suppression.

REFERENCES

[Deering89] S. Deering. "Host Extensions for IP Multicasting". Internet RFC1112, August 1989.

[Mankin98] A. Mankin, A. Romanow, S. Bradner, V. Paxson, "IETF Criteria for Evaluating Reliable Multicast Transport and Application Protocols", RFC 2357, June 1998.

[Pingali93] S. Pingali, D. Towsley, J. Kurose. "A Comparison of -Initiated and Receiver-Initiated Reliable Multicast Protocols". In Proc. INFOCOM, San Francisco, CA, October 1993.

[Floyd95] S. Floyd, V. Jacobson, S. McCanne, C. Liu, and L. Zhang. "A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing", Proc. ACM SIGCOMM, August 1995.

[Macker99] J. Macker, R. Adamson, "The Multicast Dissemination Protocol (MDP) Toolkit", Proc. IEEE MILCOM 99, October 1999.

[Holbrook99] H. Holbrook and D.R. Cheriton. "IP Multicast Channels : EXPRESS Support for Large-scale Single-Source Applications", Proc. ACM SIGCOMM, 1999.

[Levine96] B.N. Levine, J.J. Garcia-Luna-Aceves. "A Comparison of Known Classes of Reliable Multicast Protocols", Proc. International Conference on Network Protocols (ICNP-96), Columbus, Ohio, Oct 29--Nov 1, 1996.

[Nonnenmacher98] J. Nonnenmacher and E. W. Biersack, "Optimal Multicast Feedback," Proc. IEEE INFOCOMM , p. 964, March/April 1998.

[Fuhrmann01] T. Fuhrmann, J. Widmer, "On the Scaling of Feedback Algorithms for Very Large Multicast Groups", Special Issue of Computer Communications on Integrating Multicast into the Internet 24 (5-6), pp. 539-547, 2001.

[Widmer01] J. Widmer, T. Fuhrmann, "Extremum Feedback for Very Large Multicast Groups", Proc. Third International Workshop on Networked Group Communication (NGC), London, 2001.

[Abramowitz70] M. Abramowitz, I. Stegun, *Handbook of Mathematical Functions*, U.S. Department of Commerce National Bureau of Standards, Washington, DC 1970.

[ns] The Network Simulator: ns-2, <http://www.isi.edu/nsnam/ns/>.

[mdp] MDP Software Toolkit, downloadable software and documentation, <http://manimac.itd.navy.mil/MDP>.

[Adamson01] R. Adamson, J. Macker "A TCP Friendly Rate-Based Mechanism for Nack-Oriented Reliable Multicast Congestion Control," Proc. IEEE GLOBECOMM 2001, 2001.

[Handley01] J. Widmer, M. Handley, "Extending Equation-based Congestion Control to Multicast Applications", Proc. ACM SIGCOMM 2001, pp. 275-286, San Diego, CA, 2001