RELIABLE MULTICAST AND INTEGRATED PARITY RETRANSMISSION WITH CHANNEL ESTIMATION CONSIDERATIONS

Don E Gossink[1] and Joseph P Macker[2]
[1]Communications Division, Defence Science and Technology Organisation
[2]Information Technology Division, Naval Research Laboratory

ABSTRACT

This paper explores extensions to parity-based retransmission schemes applied to protocols for reliable multicast delivery. It considers a hybrid protocol scheme to potentially reduce the number of required data repair cycles through channel loss prediction. It is conjectured that such a scheme has potential benefit for particular network architectures, such as direct broadcast satellite, in improving protocol throughput delay and efficiency.

1 INTRODUCTION

An extensive body of work exists on reliable multicast networking protocols and frameworks. While there are many design aspects beyond packet retransmission, reliable multicasting approaches generally utilise packet retransmission schemes [1]. Much work has gone into improving the efficiency and scalability of packet retransmission within a multicast environment. Recently there has been growing interest and exploration in the area of applying parity-based encoding methods to reliable multicasting transport mechanisms. We focus on proposed parity-based packet retransmission schemes and potential enhancements for multicast protocols.

The integration of coded "packet healing", specifically erasure-based coding, has been shown to support improved efficiency and scalability when combined appropriately with reliable multicast transport techniques [2,3]. Our discussion builds upon previous work in this area and explores some further issues related to hybrid transport control. Work described in [2,3] demonstrated general performance gains of an integrated pure parity-based approach in supporting scalable multicast data retransmission schemes. We propose extensions to this approach by including "insurance" parity packets in the transmission process based upon predicted channel loss conditions. We also briefly explore performance issues of several multicast approaches.

2 PURE ARQ AND RELIABLE MULTICAST

When error detection is the primary means of error control, a communication system must provide a way of alerting potential transmitters that data retransmission is necessary. Standard data retransmission procedures known as automatic repeat request or automatic retransmission query (ARQ) methods are well established in practice and are generally used to achieve reliability of data transfer [6]. ARQ techniques generally involve the use of sequence numbered data packets and feedback of positive acknowledgments (ACKs) and/or negative acknowledgments (NAKs) of received packets from receivers. Most reliable multicast transport methods generally use some variant of ARQ to reliably ensure delivery of network data to a group of multicast receivers. Relative data ordering information is typically provided by sequence numbers sent as part of the transmission packet structure and lost or corrupted packets are treated by receivers as missing data objects to be replaced by successfully received data retransmission.

Several reliable multicast approaches designed for efficient bulk data transfer (e.g. multicast file transfer) additionally make use of aggregate NAK message lists or bitmaps within a transmitted information block to further reduce receiver feedback requirements [mdp, mftp, others]. In these cases, a data source collects explicit negative acknowledgments (NAKs) from multicast recipients for missing packets in a transmitted information block. Requested missing packets are retransmitted and the process is repeated until the original block is completely reconstituted at all multicast receivers.

A principal shortcoming with pure selective repeat ARQ approaches unique to multicasting is that a reduction in required retransmitted packets is accomplished by an overlap of missing data amongst receivers within the receiver group. In the case of correlated errors or losses this is true; however, for uncorrelated (independent and identically distributed) erasures the overlap is probabilistic and

0-7803-4984-9/98/$10.00 ©1998 IEEE.

different receivers will generally require re-transmission of different packets. Such a situation increases both the amount of retransmission packets and the number of receiver requests. Under this condition, as the number of receivers increases, it has been shown that the pure ARQ approaches become less efficient. The following sections review potential approaches to ameliorate these shortcomings and discuss some additional performance considerations.

3 INTEGRATED PARITY TRANSMISSION AND ARQ

Integrated parity transmission for reliable multicast transport is a technique that integrates a selective repeat ARQ approach with parity-based Forward Error Correction (FEC); related issues are described in [2,3]. The principal of operation is that in order to repair t erasures in a block, a minimum of t parity symbols need to be transmitted [4,5]. If the information coding process is done properly, the location of the packet erasures (missing data) within a block is not required information, but rather the number of lost packets is needed. This means that the transmission source can perform efficient repair if it determines the "maximum number" of erased packets amongst a multicast receiver set [3]. Multicast receivers within a network often have implicit (e.g., based on sequence numbers) or explicit (e.g., based on error detection) means of determining the status of a data stream, where packets are missing or invalid. From an information theoretic standpoint, such data can be viewed as an erasure within an information code. An erasure refers to known data location with invalid received data. As described in [2], parity data packets can be formed from blocks of information packets using well known coding techniques (e.g., Reed Solomon) to form the basis of a packet-oriented erasure code. At a particular multicast receiver, the known location of the erased packets and t received parity packets provide a sufficient condition for the original data block to be recomposed, given that t or fewer packets were originally lost.

If t' parity packets are received, where $t' \geq t$, a receiver with t erased packets will be able to reconstruct the original block of data. Herein lies much of the performance gain advantages of parity coded transmission in a multicast environment. Within a multicast group, a single parity packet has the ability to repair different missing data packets at different receivers within the multicast group. As such, within a given transmission block only the worst case performance amongst a receiver group or repairing region (e.g., in the case of localised repairing) need be considered. Upon gaining this information, data sources can prepare the proper amount of parity packets for repair transmission. In a basic adaptation of pure ARQ, receivers can inform a potential source with packet erasure information. Subsequently, the source can compile a record of the number of packets missed by the individual receivers for a given block of data. Once receivers have reported their erasure information, the "maximum reported loss" can be found and used to determine the number of group parity packets required. This would generally correspond to the worst case receiver, and the sender would prepare a block of parity packets for transmission to repair the data block. As described earlier, receivers with less than the maximum number of erased packets are being amply catered for.

It should be noted that when there are correlated errors the aforementioned process should perform no worse, on the basis of transmission efficiency, than the pure ARQ approach previously discussed. This result is also discussed in [2] and investigated in [3].

A further extension of such a protocol is possible when one assumes multicast receiver transmissions are overheard or repair reports are provided to other multicast receivers. Under this assumption, backoff reporting algorithms can be used to reduce the number of redundant receiver feedback messages. This scheme is presently implemented in existing hybrid data retransmission schemes[1]. For example, if receiver A overhears receiver B reporting n losses and A has t losses, where $t \leq n$, then receiver A can remain silent during this reporting period. Dependent upon the underlying network architecture and other considerations (e.g, unidirectional routing, timing requirements) such backoff techniques may or may not be an effective protocol optimization, but they should be considered for a general, scalable design.

It is expected that distributed, localised repairing algorithms can reduce the expected number of uncorrelated packet losses required

[1] http://manimac.itd.nrl.navy.mil/MDP

to be repaired globally within a multicast session and subsequently reduce the potential benefits of applying parity-based retransmission procedures. Such "localised repairing" techniques are not always effective or even possible: in particular networking architectures that involve large overheads, or fully-distributed timing estimates. Direct broadcast satellite, hybrid-cable access, or asymmetric wireless networks provide example scenarios where it can be prohibitive for multicast receivers to communicate directly and efficiently with other multicast receivers. In addition, even if local repairing techniques are feasible, the assumption of removing uncorrelated loss by repairing locally may not hold (e.g. as is the case when radio receiver loss is dominant). In conclusion, we see "localised repairing" and "parity repairing" as complementary schemes contributing to improving the scalability and efficiency of a wide variety reliable multicast protocol designs. After taking different architectural assumptions and criteria into consideration, it can be appropriate from an engineering perspective to apply such techniques in an integrated fashion or separately.

4 INTEGRATED RELIABLE MULTICAST FEC/ARQ WITH CHANNEL ESTIMATION

It has previously been shown that parity-based retransmission can afford considerable saving in the number of packets required for multicast group repair. In addition, the number of retransmission cycles required for reliable delivery of a complete information data block across a multicast group can be reduced. We now consider a modification to the integrated parity-based retransmission ARQ scheme. Assuming a stationary loss channel, it is apparent that additional parity in the first repair cycle can provide robustness against the corruption or loss of transmitted parity symbols. Such an approach could potentially prevent the need for multiple repair cycles and subsequently reduce the total number of retransmission cycles required to ensure the successful delivery of the overall information block to the multicast group.

Similar to the preceding discussion, maximum packet loss at any one receiver amongst a group of multicast receivers is one metric of principal interest. Using this metric, an estimate of the packet erasure probability can be made simply by dividing the number of packets lost by the total number of packets in the desired block. This estimate will be referred to herein as $\hat{\rho}$. The accuracy of $\hat{\rho}$ increases with the number of packets sent (with the assumption that channel loss is stationary). Strategies for improving the estimation of $\hat{\rho}$ for given communication channels are not discussed here and we only consider the possible use of $\hat{\rho}$ and related performance issues. Given the overall assumption that $\hat{\rho}$ provides an accurate estimate, it is considered likely that $\hat{\rho}$ times the number of repair packets will subsequently be erased across the receiver group during a repair cycle. Thus, in order to attempt to minimise the number of overall repair cycles required, it may be prudent to transmit additional "insurance" parity packets to cover anticipated packet loss. The number of repair parity packets required can be determined by:

$$\left(\frac{1}{1-\hat{\rho}}\right)\times(\text{Max No. of Packets Erased}) \qquad (1)$$

For practical implementations a ceiling function, to round the result of (1) to the nearest integer is applied.

The usefulness of the proposed technique is all the more important, because whenever possible we would like to limit the number of required repair cycles to one. This is because losses in the first parity cycle must subsequently be repaired with independent parity or with explicit loss information as in pure ARQ. It should be pointed out that in [2], the performance gain analysis only compared reduction of messaging requirements within the initial repair cycle for pure ARQ and integrated parity-based ARQ.

5 SIMULATION DESCRIPTION

The preceding discussions for the different possible reliable multicast approaches do not adequately describe the implementation requirements of a particular protocol, or the possible measures of effectiveness. Due to space considerations, a more detailed discussion of protocol implementation requirements are not provided in this paper. The following are the two measures of effectiveness for this paper: average number of packets required and average delay, K, to

reconstruct a completely error free block of information. Averages are determined from Monte-Carlo trials. The average number of required transmit packets is used to describe the amount of source transmission and retransmission packets required to successfully send a block of N packets to all M multicast users. The average cycle delay is based upon the average number of transmit cycles required to ensure reliable delivery of the information block.

The term "delay" is used in this paper to refer to the total number of transmission cycles. A cycle is the period taken from when a block of information, or repair packets, is transmitted to the time that all request for healing packets, or acknowledgment of error free replication of the information block, occurs. Thus, the average cycle delay refers to the number of cycles that on average will be experienced to achieve reliable multicast. This is not an absolute measure of time, since additional transmission within a cycle will involve additional transmission time. The actual relationship between cycle time and packet transmission time is scenario dependent. Within our discussion the first transmission cycle is the original transmission of the information block the second, third, and so on are the subsequent repair retransmissions.

The three previously discussed approaches of achieving reliable multicast have been investigated via Monte-Carlo simulation. Some common parameters of the simulations are as follows: One Sender to M receiver structure, similar to Global (Theatre) Broadcast System, assumed. Homogeneous channels for all M receivers. Packet erasures are Independent and Identically Distributed (IID). Statistically stationary channel. All receivers have indicated their repair requirements before transmission of repair packets, etc. Back link is error free. CRC is error free and identifies all errors and packets with errors are subsequently identified as erasures.

Within the simulation, random erasures are introduced by comparing random additive white Gaussian noise (AWGN) against a threshold. If the threshold is exceeded the packet is declared erased; otherwise it is assumed to arrive without incident. The determination of the threshold is done analytically using the complementary error function and a specified packet erasure probability.

6 SIMULATION RESULTS AND DISCUSSION

Simulations were conducted for various Numbers of receivers (M), Numbers of Packets in an information block (N) and various erasure rates (ρ). This was done for all three reliable multicast methods and the results are plotted on a logarithmic scale in the following figures.

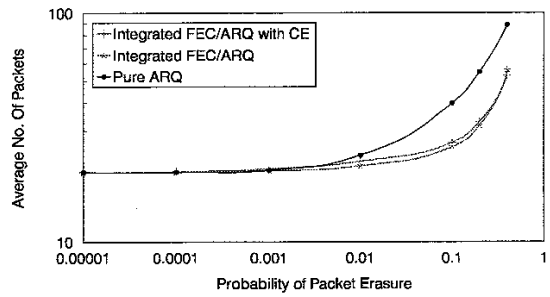

Figure 1 Average No. of packets sent versus probability of packet erasure for M=20, N=20.

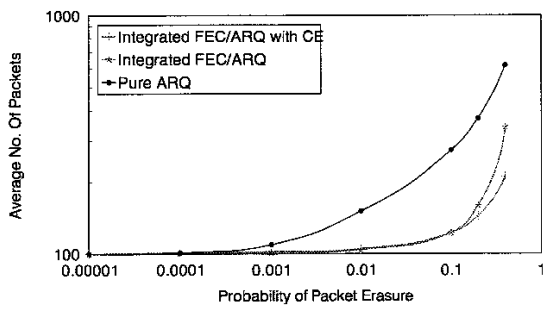

Figure 2 Average No. of packets sent versus probability of packet erasure for M=20, N=100.

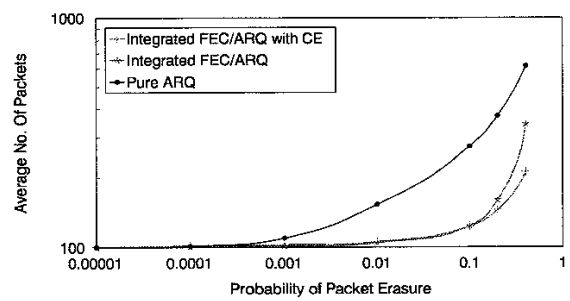

Figure 3 Average No. of packets sent versus probability of packet erasure for M=100, N=100.

The preceding plots of Figures 1 to 3 demonstrate that integrated parity-based retransmission approaches can substantially reduce the average number of transmitted packets required to ensure reliable delivery of a

complete information block across a multicast group. This indicates that the integrated parity retransmission approaches can more efficiently utilise available channel resources and improve scalability properties of existing protocols. It is interesting to note that the integrated FEC with channel estimation approach at very low probability of erasure requires fractionally more packets to be transmitted than pure ARQ and integrated FEC/ARQ. At moderate erasure probabilities the channel estimation still requires slightly more packets to be sent than integrated FEC/ARQ alone but is clearly superior to pure ARQ. At high erasure rates, the scheme using integrated FEC/ARQ and channel estimation approaches the average number of packets required of integrated FEC/ARQ. In some instances, it requires less. The aforementioned characteristic appears to be a dependent upon both *M* and *N*.

The second measure: average cycle delay may be considered more important as it indicates which approach may deliver the information in the least number of transmission cycles. This may translate to an information block being reliably received significantly sooner in time and reflect how soon sources may release buffers set aside for holding information blocks and parity data. Additionally, by avoiding the need to begin explicit individual packet retransmission requests beyond the first repair cycle we potentially reduce the workload required of the protocol under independent loss conditions. The following plots of Figures 4 to 6 indicate the simulated average cycle delay (plotted on a logarithmic scale) for the three approaches under the same assumptions as the average number of packets results.

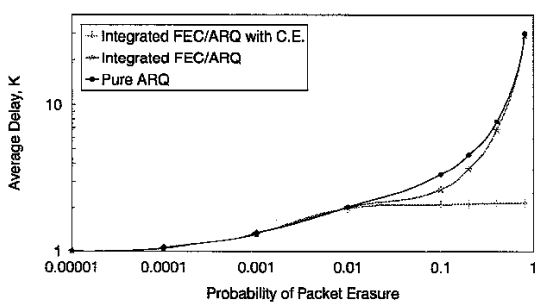

Figure 4 Average delay (No. of cycles) versus probability of packet erasure for M=20, N=20.

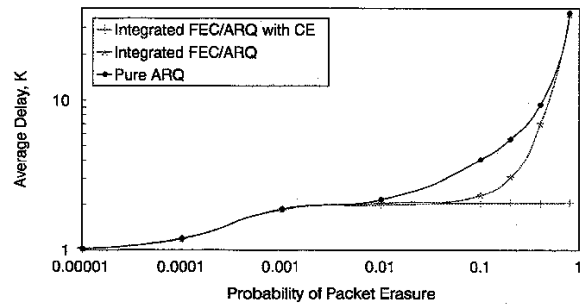

Figure 5 Average delay (No. of cycles) versus probability of packet erasure for M=20, N=100.

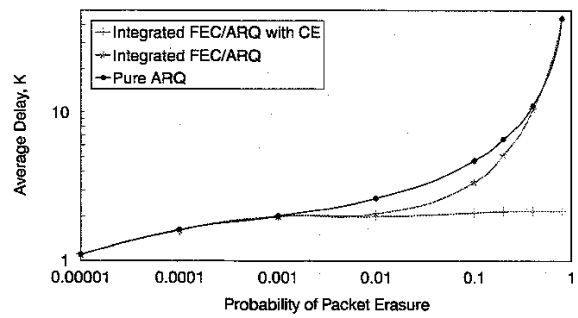

Figure 6 Average delay (No. of cycles) versus probability of packet erasure for M=100, N=100.

The plots of Figures 4 to 6 indicate that at low erasure rates the three reliable multicast approaches are approximately equivalent in terms of cycle delay as anticipated. For moderate erasure rates, both integrated FEC/ARQ approaches outperform pure ARQ with significantly lower average number of cycle delay. These performance differences become more significant as receiver group size and/or block size is increased. For high erasure rates, the integrated FEC/ARQ approaches the same number of delays as Pure ARQ. In all the plots, the integrated FEC/ARQ with channel estimation approach consistently demonstrated a lower average number of delays and when erasure rates were high on average, a fraction above two delay cycles was required. This means that the value of the "insurance" parity bits was such that the original information block and a single repair cycle of requested parity plus "insurance" parity was sufficient and reduced the likelihood of requiring additional repair cycles.

7 IMPLEMENTATION CONSIDERATIONS

The model for erasure-based repairing presented in this paper is an ideal model.

However, with practical implementations one must consider the erasure coding technique being employed. That is standard erasure coding approaches, such as Reed-Solomon, have specified rates, encoded symbols (n) and number of information symbols (k) per code set. These constraints can be massaged by lengthening and shortening [5]. Shortened codes can be formed to encode short data blocks but often involve significant processing overhead for long parity blocks. Practical implementation considerations may lead us to limit the length of parity blocks.

It can easily be shown that for any (n,k) erasure code that the integrated FEC/ARQ with channel estimation approach can be applied as long as

$$\hat{\rho} \leq \frac{(n-k)}{k} \qquad (2)$$

is satisfied. That is if the estimate ρ is larger than the constraint of (2) then one of the other reliable multicast approaches will need to be employed.

Another consideration that needs to be taken into account is the processing requirements for the reconstruction process from the parity packets. The "decoding" of code words to reconstruct the information block can require considerable processing. The decoding rate required will be dependent on the coding approach employed and the received symbol/packet rate. In some applications, this combined requirement can be considerable; restraining the amount of parity that can be afforded. With increasing processor speed or optional hardware solutions this constraint may be lessened in the future, however, using current day technology it remains an implementation concern.

8 CONCLUSIONS

In this brief study a number of aspects of reliable multicast were investigated and a novel method of including "insurance" parity within integrated FEC/ARQ was introduced and early simulation results were provided.

It was demonstrated via simulation that the integrated FEC/ARQ approaches perform better than Pure ARQ in terms of average number of packets and average cycle delay required to achieve reliable multicast delivery. It was also demonstrated that integrated FEC/ARQ with channel estimation offers potential improvements, especially when delay is an important issue. The use of "insurance" parity packets, with accurate loss estimators, is particularly interesting as it appears to approach bounded cycle delay over an extended range of erasure rates. This result is anticipated to provide significant benefit in network architectures with significant delay characteristics (eg., asymmetric satellite dissemination). While the results and analysis here are largely preliminary these results should be considered in future design of parity-based ARQ schemes, especially when the statistical packet loss estimation is seen as significant and sufficiently stable within the repairing time window.

9. REFERENCES

[1] Yamamoto, M., Kurose J F., Towsley D F. and Ikeda H., "A Delay Analysis of Sender-Initiated and Receiver-Initiated Reliable Multicast Protocols", Infocom 1997.

[2] Macker, J P., "Reliable Multicast Transport and Integrated Erasure-Based Forward Error Correction", In Proceedings of MILCOM 97.

[3] J. Nonnenmacher, E. Biersack, D. Towsley. "Parity-Based Loss Recovery for Reliable Multicast Transmission," Technical Report 97-17, Dept. of Computer Science, Univ. of Massachusetts, March 1997.

[4] J.G. Proakis, *Digital Communications (Second Edition)*, McGraw Hill International Editions Computer Science Series, Singapore, 1989.

[5] S.B. Wicker, *Error Control Systems for Digital Communication and Storage*, Prentice-Hall, Engelwood Cliffs, 1995.

[6] B. Sklar, *Digital Communications: Fundamentals and Applications*, Prentice Hall,1988.