

TECHNICAL RESEARCH REPORT

Reliable Multicasting via Satellite: Delay Considerations

by Stephen M. Payne and John S. Baras

CSHCN T.R. 2000-10
(ISR T.R. 2000-28)



The Center for Satellite and Hybrid Communication Networks is a NASA-sponsored Commercial Space Center also supported by the Department of Defense (DOD), industry, the State of Maryland, the University of Maryland and the Institute for Systems Research. This document is a technical report in the CSHCN series originating at the University of Maryland.

Web site <http://www.isr.umd.edu/CSHCN/>

RELIABLE MULTICASTING VIA SATELLITE: DELAY CONSIDERATIONS*

Stephen M. Payne and John S. Baras

Center for Satellite and Hybrid Communication Networks
Institute for Systems Research
University of Maryland
College Park, MD 20742

ABSTRACT

Many different reliable multicast protocols have been proposed and analyzed in the current literature. Since satellites are naturally a broadcast medium, multicast communications have the potential to greatly benefit from their wide-scale deployment. The performance of reliable multicast protocols needs to be studied and become better understood over networks including satellite links. Most of the analysis performed on these protocols has dealt with bandwidth usage, buffer requirements, and processing delay. Very few studies address the transmission delay incurred from using reliable multicast protocols. Hybrid error control protocols have been studied in terms of bandwidth and delay. The effects of different estimation schemes coupled with autoparity usage are investigated and results are compared. Simple adaptive mechanisms used with a local recovery scheme are found to offer the best overall results in terms of reducing recovery latency and satellite bandwidth usage.

INTRODUCTION

Multicast communications are important for hierarchical wireless networks, such as those used in the military. Ensuring reliable transmission in multicast communication is key for military networks. Terrestrial, wireless, and mobile LANs supported by satellite provide great promise for reliable multicast communication. This paper discusses the combination of existing schemes required to provide reliable multicast communications in military networks.

Many different reliable multicast protocols have been proposed and analyzed in the current literature. Examples of such protocols include RMTP[8], Reliable Adaptive Multicast Protocol (RAMP) and Multicast File Transfer Protocol (MFTP). RAMP was intended for use in collaborative military applications such as simulated war games.

*Prepared through collaborative participation in the Advanced Telecommunications & Information Distribution Research Program (ATIRP) Consortium sponsored by the U.S. Army Research Laboratory under the Federated Laboratory Program, Cooperative Agreement DAAL01-96-2-0002. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

It attempts to reliably deliver multicast data while reducing latency [2]. MFTP [4] was designed for the reliable non-real-time bulk transfer of data. Since latency is not a critical design constraint, this protocol sacrifices delay to gain extra scalability and universal operation over different network infrastructures including satellite and other asymmetric environments.

As seen from the above examples, reliable multicast protocols have been designed for specific applications in specific network environments. As reliable multicast applications begin to require low latency operation over hybrid networks, reliable multicast protocols need to be studied in such networks. A natural starting point for such studies is the consideration of the delay characteristics over satellite links. Satellite links suffer from relatively high raw bit error rates compared with terrestrial fiber links as well as higher delay characteristics. Therefore, with regards to reliable multicast applications, satellite communication provides an interesting set of technical complications in which latency becomes an important performance metric.

When considering latency of reliable multicast applications over satellite links, error recovery becomes a crucial issue. The use of an ARQ scheme requires a feedback channel for proper protocol function. However, due to the high latency over satellite links, such schemes experience significant throughput degradation (e.g. the TCP degradation observed over high latency links [1]). For this reason, as well as the relatively large amount of feedback bandwidth needed for multicast ARQ implementations, one realizes the importance of applying packet level Forward Error Correction (FEC). As previously demonstrated in [5], hybrid error control (HEC) schemes that use parity packets to reduce the residual packet error probability and a feedback mechanism to ensure reliability combines the advantages of both FEC and ARQ schemes to form a more ro-

bust protocol. Variations of these HEC protocols offer additional potential benefits. One important variation uses local recovery between terrestrially connected nodes. This variation limits the amount feedback to and the number of transmissions carried over the satellite link.

The remainder of the paper is organized into four sections. The next section contains the descriptions of the generic protocols and the two network scenarios considered in this paper. These two network scenarios are then studied and their results are presented. The paper concludes by summarizing results and presenting areas for future work

GENERIC PROTOCOL & NETWORK SCENARIOS

The generic HEC protocol shown in Figure 1 is similar to the one studied in [5] and is assumed throughout this paper.

- 1)The sender sends a transmission group of k data packets and a (*autoparity*) ϵh parity packets from the associated FEC block
- 2)All packets within the transmission group (TG) can be recovered if there are fewer than a missing packets among the $k + a$ transmitted packets.
- 3)During the initial transmission round, a receiver detecting more than a missing packets requests the number of parity packets required to complete the TG. In subsequent retransmission rounds, the receiver requests the number of packets required to complete the TG.
- 4)The sender multicasts the maximum number of requested parity packets from all receivers until all parity packets associated with the TG have been used. At that time packets requiring retransmission are placed into a new transmission group.

Figure 1: Generic HEC protocol

The parity packets sent during each retransmission round are subject to the same error probability as the data packets. Lost parity results in additional retransmission rounds (step 4:Figure 1) to ensure reliable delivery. As parity was sent in the initial round, extra or “insurance” parity can also be sent in subsequent retransmission rounds. This concept is referred to as channel estimation considerations. For example, the amount of parity sent can be calculated using the maximal packet loss probability as measured during the initial transmission round (step 3:Figure 1) [3].

There are two basic network scenarios that are considered in this paper; the unconnected cluster scenario and the

connected cluster scenario. A connectivity cluster is defined as set of nodes that are virtually connected with each other. In each cluster, it is assumed that only one node called a Privileged Receiver (PR) can communicate with the satellite. This assumption requires that all other nodes within a particular cluster be connected (either directly or indirectly) to their corresponding PR. These non-privileged nodes are known as Dependent Nodes (DN).




Figure 2: Unconnected Cluster Scenario

In the unconnected scenario as seen in Figure 1, a source sends multicast data over a satellite to R connectivity clusters. When considering reliable multicast to a group consisting of both PRs and DNs, a hierarchical approach immediately presents itself as a viable option. There have been studies (e.g. [6], [10]) that indirectly purport the use of hierarchical multicast to reliably disseminate data. As all packets destined for the connectivity cluster must pass through the PR, each PR has the opportunity to buffer these packets. If the PR correctly receives k packets (either data or parity packets), then it reconstructs the original k data packets. If the PRs have the ability to create new parity packets, then these parity packets combined with the original data packets can be used to locally satisfy retransmission. Such a scheme is similar to the APES scheme, SDBR proposed in [10].

Due to the fact that PRs cannot generate parity unless they have the k original data packets, reliable delivery can be subdivided into two stages; 1) delivery from source to the PRs, and 2) delivery from PRs to nodes within connectivity clusters. Considering only the first stage reduces the problem to the reliable delivery of k data packets to R receivers. It is assumed that the PRs do not forward packets until they correctly receive and decode the k data packets.

Upon the successful reception of the k data packets, each PR assumes the responsibility as the multicast source for its corresponding connectivity region. More detailed explanation of different possible delivery schemes are presented in [9].

The connected scenario differs from the previous scenario in that the connectivity clusters overlap (see Figure 3). Since the connectivity clusters are terrestrially connected, a multicast tree can be established between the PRs. In such a scheme, the source sends the initial transmission to all PRs within a specified multicast group. When a PR detects a loss, it starts a local retransmission cycle by setting a local retransmission timer. If no members in the local neighborhood respond prior to the timer's expiration, then the PR begins a global recovery phase in which it transmits a global NAK over the satellite to the source. The PR enters another local recovery cycle if the requested packet is not received. This process is repeated until the entire TG is reliably delivered to all PRs of the multicast group.




Figure 3: Connected Cluster Scenario

Now, the assumptions concerning the packet error probability at receivers in the multicast group need to be considered. Packet errors (packet losses) are assumed to be both spatially and temporally independent at the receivers. Independent losses uniformly scatter the losses amongst the receivers in the group; whereas shared losses have the potential to concentrate losses in particular areas of the tree. In reality, the errors occurring at receivers depend upon many different factors. When using tree structures, a loss within the tree will be shared by more than one receiver (i.e. shared loss). Multicast trees' shared losses are modeled well by a full binary tree (FBT) [7]. Regardless of the locality of the losses, the source still needs to trans-

mit parity repairs over the entire tree. For a more detailed explanation of the impact of shared losses and temporal losses refer to [9].

The results for the infinite parity cases are presented in the remainder of this paper. Using a finite number of parity packets negatively impacts performance. The degree of this impact depends upon the actual number of parity packets generated [9]. In the unconnected scenario, several simplistic adaptive mechanisms are studied. In the connected case, three local recovery schemes with different local versus satellite transmission round ratios (LxS) are studied. Only the original source is able to use autoparity and channel estimation techniques during satellite (re)transmission rounds. All local requests are fulfilled using distinct parity packets.

RESULTS: UNCONNECTED CLUSTER SCENARIO

A large amount of analytical work was completed in [9]. In addition to this work, simulations were created so that results could be obtained in the cases where analytical results were not tractable in addition to verifying the analytical work. The number of simulation trials used to create each data point were no less than 1,000. Although not extensively studied, it was found that increasing the number of trials to 10,000 resulted in no appreciable difference in results. Simulations were created for each of the four possible variations on the unconnected connectivity cluster scenario; infinite parity without channel estimation considerations, infinite parity with channel estimation considerations, finite parity without channel estimation considerations, finite parity without channel estimation considerations. The finite parity variations are not discussed in this paper.

In the unconnected scenario, the use of autoparity was examined and found to reduce the number of transmission rounds (as shown in Figures 4 and 5). This reduction comes at the cost of additional bandwidth usage for low packet loss probabilities.




Figure 4: $E[\text{Transmission Rounds}]$ for $a=1$

However, as the packet loss probability increases, the intrinsic overhead cost associated with autoparity is reduced (see Figure 6). Such a reduction suggests that intelligent use of autoparity can help to lower the delay without drastically increasing bandwidth usage. Secondly, the maximal packet loss probability, channel estimation technique [3] was investigated using different levels of autoparity. It was found that the number of required retransmission rounds was reduced to approximately two at the cost of a substantial increase in the number of packets.




Figure 5: $E[\text{Transmission Rounds}]$ for $a=3$




Figure 6: Ratio of expected values

Before studying the connected scenario with local recovery schemes, simplistic adaptive mechanisms where the parity provided during each transmission is adjusted based upon observed packet loss statistics are investigated. One possible scheme dynamically changes the amount of autoparity based upon a moving average of the packet loss probability observation. The performance of this dynamic autoparity technique does not outperform the $a=3$ case over the entire range of packet loss probabilities. However, this scheme is more bandwidth efficient at lower packet loss probabilities and does not drastically increase the expected number of retransmission rounds. Another possible improvement uses a moving average estimation of the packet loss probability to calculate the number of “insurance” packets that are sent during subsequent retransmission rounds. This moving average scheme apply to insurance packets does not perform as well as the maximal channel estimation technique at higher packet loss probabilities. [9]

These two techniques can be combined into one protocol. When comparing the dynamic autoparity case protocols that use and do not use the moving average channel estimation technique, one notices that the combination of these two scheme results in a decrease in the number of retransmission rounds (see Figure 7 and Figure 8). In both of these figures the TG size, k , equals 20. This performance improvement is realized with a negligible increase in the expected number of transmitted packets [9].




Figure 7: $E[\text{Transmission Rounds}]$ for dynamic autoparity




Figure 8: $E[\text{Transmission Rounds}]$ for dynamic autoparity with the moving average channel estimation technique

RESULTS: CONNECTED CLUSTER SCENARIO

For the connected cluster scenario, there are several alternative solutions that could limit the expected number of satellite rounds. Specifically, local network usage was examined and found to reduce both the expected number of satellite transmissions and the expected number of satellite transmission rounds. The first local recovery mechanism studied alternated between one global and one local transmission round. This scheme was examined in both the absence of channel estimation techniques and in their presence. Without channel estimation, the local recovery schemes used fewer satellite rounds and transmitted fewer satellite packets than their non-local recovery counterparts. In the presence of the maximal channel estimation technique, the $1x1$ local recovery scheme did not offer substantial gains over the non-local recovery case. When increasing the LxS ratio (e.g. from $1x1$ to $3x1$), the performance gains depended upon the number of PRs within the

terrestrial multicast tree as well as the number of DNs within each connectivity cluster. For smaller connectivity clusters, there are fewer potential DNs to correctly receive the entire TG. For larger ones, there are more local subtrees and therefore a greater chance that additional satellite rounds are required. More in depth discussion and accompanying graphs can be found in [9].

LxS local recovery schemes coupled with dynamic autoparity and the moving average channel estimation scheme offer noticeable performance gains over their non-local counterparts. In terms of both bandwidth and delay, these combination LxS local recovery schemes offer the lowest average number of retransmission rounds using the fewest number of packets. In fact, using a sufficiently high LxS enables the protocol to decrease the number of satellite transmission rounds to approximately less 1.15 for multicast groups larger than four receivers [9]. For large groups, this is a significant improvement over the previously discussed results.

CONCLUSIONS & FUTURE WORK

Although most of the results obtained in this paper were for the expected number of satellite transmission rounds, this metric gives a good feel for the amount of time the protocol requires to reliably deliver a TG. By using relatively simple adaptive techniques, the “combination” protocols were found to decrease the expected number of rounds. In the connected scenario case, performance improvements come at the cost of local transmission rounds and therefore places a larger load on the local networks. Even though this study demonstrated the performance enhancements achieved through using relatively simple modifications to existing reliable multicast protocols, further efforts should be applied to the following areas. More realistic local packet loss probabilities that include correlated losses need to be incorporated into study. The preceding study assumed that all receivers have the same processing power. If they do not have similar capabilities, then the schemes suggested in this paper need to be adapted to account for this additional design constraint. The local recovery scheme in which both local parity packets and original data packets are transmitted as repairs can be studied. Perhaps, the most promising area of further research lies in obtaining and studying satellite packet loss statistics and developing advanced estimation schemes.

REFERENCES

- [1] M. Allman et al., "Ongoing TCP Research Related to Satellites." draft-ietf-tcpsat-res-issues-09.txt.
- [2] R. Braudes, S. Zabele. "Requirements for Multicast Protocols" RFC 1458, May 1993.
- [3] D. Gossink, J. Macker. "Reliable Multicast and Integrated Parity Retransmission with Channel Estimation Considerations". IEEE Globecom, 1998.
- [4] K. Miller, K. Robertson, A. Tweedly, M. White. "StarBurst Multicast File Transfer Protocol (MFTP) Specification". Internet Draft, Work in Progress, draft-miller-mftp-spec03.txt, April 1998.
- [5] J. Nonnenmacher, E. Biersack, and D. Towsley. "Parity-based loss recovery for reliable multicast transmission". IEEE/ACM Transactions on Networking, vol.6, no.4, p. 349-61, August 1998.
- [6] J. Nonnenmacher, M. Lacher, M. Jung, E. Biersack, G. Carle. "How bad is reliable multicast without local recovery?". Proc. IEEE INFOCOM '98, vol. 3, p. 972-9, 1998.
- [7] J. Nonnenmacher and E. W. Biersack. "Performance modelling of reliable multicast transmission", Proceedings IEEE INFOCOM '97, vol. 2, p. 471-9, April 1997
- [8] S. Paul, K.K. Sabnani, J.C. Lin, and S. Bhattacharya, "Reliable multicast transport protocol (RMTP)". IEEE Journal on Selected Areas in Communications, vol.15, no.3, p. 407-21, April 1997.
- [9] S. Payne. "Reliable Multicast via Satellite". M.S. Thesis, CSHCN Technical Report MS 99-4 ISR Technical Report MS 99-7, University of Maryland. August 1999.
- [10] D Rubenstein, S. Kasera, D. Towsley, and Jim Kurose. "Improving Reliable Multicast Using Active Parity Encoding Services (APES)". Proceedings IEEE INFOCOM 1999 vol. 3, p. 1248-1255, March 1999.

*The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied of the Army Research Laboratory or the U.S. Government