

# Failure Handling in a Reliable Multicast Protocol for Improving Buffer Utilization and Accommodating Heterogeneous Receivers

Gunjan Khanna, Saurabh Bagchi

*Dependable Computing Systems Lab  
School of Electrical & Computer Engineering,  
Purdue University  
West Lafayette, IN.  
Email: {gkhanna, sbagchi}@purdue.edu, john.rogers@bankofamerica.com*

John Rogers

*Architecture, Standards & Security Eng.  
Bank of America Corporate  
Information Security  
Charlotte, NC.*

## Abstract

*Reliable multicast protocols are an important class of protocols for reliably disseminating information from a sender to multiple receivers in the face of node and link failures. A Tree-based Reliable Multicast Protocol (TRAM) provides scalable reliable multicast by grouping receivers in hierarchical repair groups and using a selective acknowledgment mechanism. In this paper, we present an improvement to TRAM to minimize the resource utilization at intermediate hosts and to localize the effect of slow or malicious receivers on normal receivers. We present an evaluation of TRAM and TRAM++ on a campus-wide WAN without errors and with message errors. The evaluation brings out that, given a constraint on the buffer availability at intermediate hosts, TRAM++ can tolerate the constraint at the expense of increasing the end-to-end latency for the normal receivers by only 3.2% compared to TRAM in error-free cases. When slow or faulty receivers are present, TRAM++ is able to provide the same uninterrupted quality of service to the normal nodes while localizing the effect of the faulty ones without incurring any additional memory overhead.*

**Keywords:** *tree based reliable multicast, buffer utilization, message errors, slow or unresponsive receivers, end-to-end latency.*

## 1. Introduction

The ability to reliably transmit and update large amounts of data and content in real-time is crucial in a large number of domains, such as for financial services, content delivery network providers, large retail chains, application service providers (ASPs), interactive TV, wireless entertainment, and e-learning providers. It can take considerable time and bandwidth if the sender must send a separate copy to each receiver. IP multicasting allows a sender to distribute data to all interested parties while minimizing the use of network resources. Many applications, however, require reliable data delivery which

is resilient to failures of nodes and network links. This requirement is supported by reliable multicast protocols.

A Tree-based Reliable Multicast Protocol (TRAM) is designed to provide multicast reliability that scales to a large receiver population. TRAM ensures reliability by using a selective acknowledgment mechanism and scalability by adopting a hierarchical tree-based repair mechanism. The receivers and the sender of a multicast session dynamically form repair groups. These repair groups are linked together hierarchically to form a tree with the sender at the root of the tree, the receivers at the leaves and entities called Repair Heads (RH) at the intermediate levels of the tree. The RHs cache the messages that pass through them and initiate local repair by resending messages that are nacked by the receivers.

From an end-to-end perspective, we consider latency and jitter as parameters of interest for receivers of a multicast stream. We consider the scenario of a video stream being delivered by TRAM and measure its output parameter values. In addition to the performance metrics, there are resource utilization metrics which characterize a reliable multicast protocol. TRAM ensures the recovery from missing messages by nacks and retransmissions by the RHs. The RHs therefore need stable storage to buffer messages that have not yet been received by all the receivers. The aggregate buffer utilization in the system, combined over the sender and all the RHs, is taken as one indicator of resource usage in the system. In this paper, we evaluate the existing TRAM implementation with respect to performance and resource utilization, with varying number of receivers. The evaluation is performed on a campus-wide wide area network (WAN) with receivers located at different distances from the source. The ability of the reliable multicast protocol to tolerate faults is evaluated under different message error conditions. The errors injected are message drops, reorders, and delays, and the performance metrics are measured under these conditions.

It has to be considered that in a large group of multicast receivers, not all the receivers will be identical. Considering a practical scenario where different receivers are present in different subnets, each has varying traffic

conditions and processing constraints. Some of the receivers will be slow in processing the messages and generating acks. Some receivers may be unresponsive for long periods of time, either because of a natural failure or performance bottleneck, or because of malicious purpose. The design point should be that the QoS measures of the normal receivers should be affected as little as possible because of the slow and unresponsive receivers. We propose a protocol based on TRAM, called *TRAM++*, which localizes the effect of slow or unresponsive receivers on the correctly functioning receivers. *TRAM++* provides a mechanism to track unresponsive receivers and beyond a threshold, prune the receivers. Since the receivers are not homogenous, maintenance of synchrony among the receivers is sacrificed as a deliberate design choice to provide fairness to normal receivers. Pruning is considered a necessary step to keep the system performance above a tolerable limit, bound the resource usage and to eliminate malicious receivers.

*TRAM++* also optimizes the aggregate buffer usage in the system under normal conditions without any degradation of the performance parameters. It achieves this by dropping the guarantee of recovery of missing messages from the immediate RH to which the receiver is connected. The RH may drop messages from its buffer in order to accommodate new messages. The guarantee of recovery is provided by the sender which buffers messages in stable storage till they have been acknowledged by every receiver.

The study shows that both TRAM and *TRAM++* scale well with respect to latency as the number of receivers is increased, up to a total of 30. TRAM is not successful in isolating the normal receivers from the effect of faulty or malicious ones. *TRAM++* is able to achieve this through its protocol of differentiated acks and buffer management. *TRAM++* under a constraint of 16% of the TRAM buffer availability at the RH is able to maintain the latency within an overhead of 3.2% in the error free scenario. It is also able to prune malicious receivers faster because of local decision making ability at the RH without any additional memory overhead. It is also able to isolate and prune a malicious receiver that constantly joins and leaves the multicast group in an attempt to cause denial of service in the system. The rest of the paper is organized as follows. Section 2 presents previous related work. In section 3, we present the TRAM and the *TRAM++* protocols. Section 4 details the implementation of *TRAM++*. Section 5 gives the experiments and results, under no errors and with errors. Section 6 concludes the paper.

## 2. Related Work

IP multicast is a protocol that defines a mechanism for one or more senders to send data to a group of receivers—a collection of one or more hosts identified by

a single class D IP address [5]. The abstraction of multiple hosts into a single entity provides several advantages over unicast, or "point-to-point," protocols: the sender(s) only manages a connection with the group, not all of its members; and data can be sent to all the receivers without multiple transmissions from the sender(s). In situations where similar data is sent to multiple hosts, multicast IP can provide bandwidth conservation, reduced resource requirements for the server, and an increase in overall scalability.

Though the general idea of IP multicast is simple, the message delivery is purely best-effort and there is no guarantee that the data is going to reach all the multicast receivers. The loss of messages may be due to congestion, failures of nodes, failures of links, etc. To address these concerns, an enhanced category of multicast IP called reliable multicast has evolved. It defines mechanisms to monitor host population, congestion, transfer speeds, and data loss on a per-host basis, and mechanisms to counter the loss, thus providing a more reliable multicast environment.

Due to the wide variety of applications which require reliable multicast, it is considered that a "one size fits all" protocol is infeasible. Therefore, three classes of protocols have been proposed in the literature – (i) NACK oriented protocols; (ii) Tree-based ACK oriented protocols; (iii) Asynchronous layered coding protocols that use forward error correction (FEC). Of these classes, the second or a hybrid of the first two classes is of relevance to our current work. Several protocols have been presented and studied [8,12,13]. We chose TRAM as the representative protocol for study since it had openly available source code and a large active user community which was extremely useful in setting up and troubleshooting the system. The detailed evaluation of TRAM presented here is expected to shed light on the other protocols too since many design decisions are similar, such as local recovery at the repair head (or the counterpart in the particular protocol), and ack aggregation. The concern about constraining resource usage at intermediate hosts and localizing the effect of slow or malicious receivers applies to all reliable multicast frameworks.

TRAM was first presented by Chiu *et al* in [2]. This study evaluated some of the parameters considered here, namely, rate, loss and cache occupancy, but using a simulation model. The simulation model made several simplifications – all RHs were at a fixed distance from the sender, no node failures were simulated and only a small subset of the links (between the RH and the receivers) were injected with failures. The congestion control mechanism in TRAM was studied in a recent paper by Chiu *et al* [4]. The two aspects of congestion control – receiver feedback based windowing and server data rate based traffic shaper – are studied using an implementation for a LAN and an emulator for a WAN. The study showed

how to dynamically adjust the data rate used to schedule packet transmission at the sender to smooth the transmission. The authors in [3] examine the issue of pruning decisions in multicast transport protocols. The decision boils down to choosing a minimum data rate and pruning receivers that fail to meet the minimum rate. It is mentioned that the minimum rate has to be chosen carefully so as not to prune genuine receivers experiencing occasional network bottlenecks. Determining an optimal data rate for the system is dependent on the kind of network and its traffic conditions, and is a complex decision. The rate of the entire group is controlled by the slowest receiver in the unpruned set of receivers. This results in slowdown observed by normal receivers even in completely disjoint parts of the multicast tree. The scheme in [3] is a simplification of the more general scheme called *optimal pruning* described in [7]. In that paper, the authors propose multiple subgroups of receivers with a utility function for individual nodes in the subgroup and one for the subgroup as a whole. The algorithm presented aims to pick the subgroup that maximizes the sum of node utility and subgroup utility. We believe that the algorithm can work if it is possible to assign utility functions for all possible subgroups. In the practical scenario considered in our paper, it was not possible to come up with a utility assignment. Also, the pruning decision needs to be rapid to isolate malicious receivers and the choice algorithm in Jiang's work runs in exponential time.

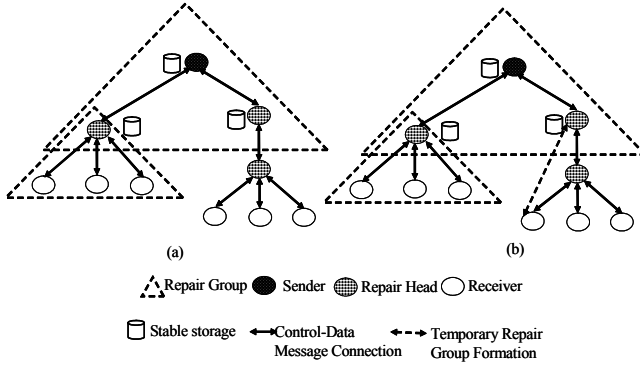
The design point of providing stable storage only at end points has been proposed and implemented in the context of publish-subscribe systems for a system called Gryphon to provide reliable exactly-once message delivery in the face of node failures [1]. The work assumed the extreme design point of no stable storage available at the intermediate nodes whereas in our study this is a parameter that can be tuned.

A promising approach to building reliable distributed systems is group communication [6]. Reliable group communication is an important paradigm to build distributed applications that require multi-peer interaction with different levels of consistency. At least a decade of research has gone into designing and improving group communication protocols to provide membership, multicast and ordering services. However, such protocols are overkill for our goal. First, the members of a group are considered homogeneous and group joins and leaves are made visible to all the group members. Second, the design principle in our system is not to enforce synchronicity among the group members which is what the group communication protocols are designed to achieve. Finally, the group communication paradigm is particularly suited to local area networks with tight bounds on end-to-end latency which may not be achievable in the wide area network considered

## 3. Protocol Description

### 3.1. TRAM

The detailed description of TRAM can be found in [2]. We provide an overview here and present details of the features relevant to the study. TRAM is distributed as a part of the Java Reliable Multicast Service (JRMS) by Sun Microsystems [11]. JRMS is a set of libraries and services for building multicast-aware applications. TRAM is designed for high scalability targeted towards multicasting streaming data from a single sender to a large number of receivers. TRAM ensures reliability by using a selective acknowledgement mechanism. An ack is sent in the form of an offset and a bit vector once every ack window (32 packets). It provides scalability by adopting a hierarchical tree-based repair mechanism. The receivers and the data source of a multicast session in TRAM interact with each other to dynamically form repair groups. These repair groups are linked together in a hierarchical manner to form a tree with the sender at the root of the tree. Figure 1 shows a typical TRAM repair tree. The nodes participating in TRAM play three roles, some nodes playing multiple roles – sender, receiver and repair head (RH). Every repair group has a receiver that functions as a group head; the rest function as group members which are said to be affiliated with their head. All members receive data multicast by the sender. The group members report lost and successfully received messages to the group head using a selective acknowledgement mechanism. Every ack message contains a start message number indicating the first missing message, and a bit vector, with a 1 denoting a missing packet and a 0 denoting a received packet. If no packets are missing, the message number indicates all messages prior to and including this one has been received and the bit vector is of zero length. An ack message is sent after every *ack window* worth of packets has been received, or an *ack interval* timer goes off. The RHs cache every message sent by the sender and provide repair service for messages that are reported as lost by the members. The RH's maintain a high and low water mark for monitoring cache occupancy. If the amount of buffer occupied by the packets goes beyond the high water mark, an attempt is made to purge the cache. Failure to do so is taken as an indication of congestion in the network. The RHs aggregate acks from all its members and send an aggregate ack up to the sender to avoid the problem of ack implosion. The data rate sent out by the sender is bounded by maximum and minimum rates configured at the sender. Receivers that cannot keep up with the minimum data rate can be pruned from the repair tree.



**Figure 1. (a) TRAM hierarchical repair tree (from [2]) (b) TRAM++ hierarchical repair tree with temporary reaffiliation**

Figure 1(a) shows a TRAM deployment with a sender, two levels of RHs and multiple receivers connected through links over which bi-directional data and ack messages flow. Two examples of repair groups are shown, one involving the sender and the two RHs at the first level, and the second showing a RH and its receivers.

### 3.2. TRAM++

TRAM++ maintains the hierarchical structure of TRAM. It builds upon TRAM with the following two goals:

1. Reduce the resource requirement at the repair heads, chiefly cache utilization, but also processing.
2. Sandboxing the slow or malicious receivers and prevent the correctly functioning receivers from being affected.

To achieve these goals, TRAM++ introduces the changes described below.

**Buffer management at RH:** The design point in TRAM++ is that the RHs may be spread over a wide area and have constraints on available buffer, while the sender has higher, though not infinite, buffer capacity. TRAM++ optimizes the buffer requirement at the RHs by pruning old messages even if they have not been acknowledged by all its receivers. The advantage is that this frees up the buffer resources at the RH for accommodating new messages which are required for the well-behaved receivers to make progress. Consequently, a nack from a receiver may not always be satisfied locally at the immediate RH. A message is not discarded from the sender's storage till it has been acked by *all* the receivers. Therefore, a nack can always be satisfied by the sender. When a RH cannot satisfy a nack, it indicates to the receiver to initiate a *temporary re-affiliation* with a RH at a higher level. A figure of the hierarchical structure in TRAM++ with a sender, receivers and repair heads is shown in Figure 1(b). Temporary re-affiliation is shown through the dotted arrow, where the receiver re-affiliates temporarily with a higher level RH for recovering the messages its RH does not have. The reaffiliation is

transient and lasts for the duration of recovery of the single packet. This process is repeated recursively if recovery is not successful at the higher level, till the receiver reaffiliates with the sender at which point its nack is guaranteed to be satisfied.

**Handling Slow or Malicious Receivers:** TRAM lets the data rate be driven by the slowest receiver. Therefore, the effect of a slow receiver is visible to the correctly functioning receivers all across the network. Even if the feature of TRAM that allows slow receivers to be pruned is turned on (which it is not for most deployments), the threshold minimum tolerable data rate is set quite conservatively. Hence, there are likely to be large periods of slowdown to the normal receivers. On the contrary, TRAM++ localizes the disruption to the part of the tree where the lagging receivers reside. In TRAM++, the RH uses two types of acks – a *greedy ack* and a *permanent ack*. The maximum sequence number of the packet that the sender should send down is sent piggybacked with acks. The greedy ack is sent by the RH upwards when its buffer reaches the low water mark. The purpose of sending the greedy ack is to indicate to the sender to send new data down even though all the receivers have not acked yet. The permanent ack is sent once all the receivers have acked. The role of this ack is to let the sender know that reclamation of storage is possible. In TRAM, the sequence number sent upwards is determined by the slowest receiver thus affecting the data rate observed by all the receivers, normal or laggard. In TRAM++, the sequence number is determined by the RH's available buffer capacity. Incorporating the additional ack requires additional computation at the sender and the RHs, which is the same cost as for the basic ack determination in TRAM. But in a failure free system where all receivers are keeping up with the data rate, the RH's buffer space is not put under pressure and therefore, the greedy acks are not sent and the additional processing overhead is not incurred.

TRAM++ has the functionality to prune receivers which are considered lagging beyond an acceptable degree. The metric used for the pruning decision is the percentage of retransmission requests which cannot be locally satisfied and require temporary re-affiliation. The percentage is calculated as a fraction of the total number of packets. When the metric exceeds a tunable threshold parameter, the receiver is pruned. This is an effective means of removing malicious receivers which may increase the processing load in the system by requesting repeated retransmissions. Alternately, this may serve as an indication to the receiver to disaffiliate from the current RH because of its resource constraints and reaffiliate with a more resource rich RH.

## 4. Implementation

In this section, we present the details of the TRAM implementation which are of importance to the protocol and the changes that have been made in TRAM++.

### 4.1. TRAM Implementation

The TRAM code is multi-threaded. These threads are responsible for carrying out the group management functions in addition to basic sending and receiving of data packets. *GroupMgmtThread* is the main thread which is responsible for starting up TRAM, initiation of the *beacon messages* by the sender and affiliation of the receivers to the senders or repair heads. The beacon messages are used to advertise the session and invite nodes to join the multicast session. This thread performs the task of sending periodic hello messages among the receivers and its head. Each receiver also maintains a backup list of heads which it can switch to if the current head resigns or fails. Once the data transmission phase starts, *InputDispThread* and *OutputDispThread* come into picture. *OutputDispThread* transmits the packets. *InputdispThread* gives the packet to all the listeners and then, each entity calls its received packet method to get the desired packet. Figure 2 shows pictorially the threads or methods which are used for upstream and downstream communication in TRAM. In the case of errors, the downstream path is identical to the error free case. In the upstream path, the receiver sends nacks to the RH which transmits the requested packets; the RH adjusts the data rate and sends the adjusted rate to the sender which finally adjusts its own rate.

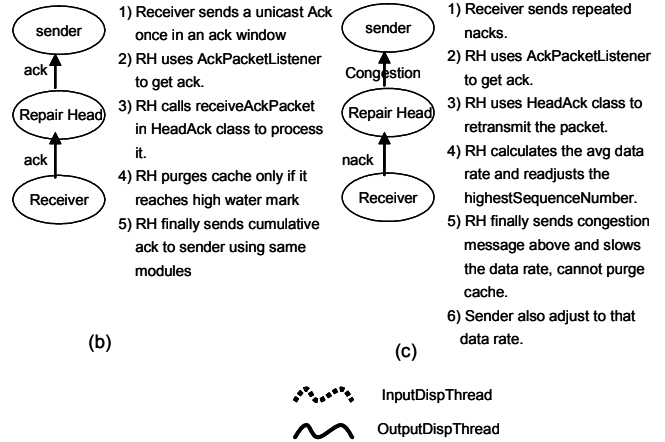
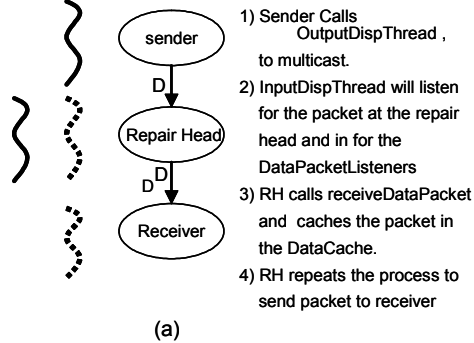


Figure 2. TRAM message processing (a) downstream with no errors, (b) upstream with no errors, (c) upstream with message, node or link errors

### 4.2. Modifications for TRAM++

To create TRAM++ from TRAM, we have added new message types to the existing message and sub-message types of TRAM. We have tried to minimize the changes to the basic TRAM structure and be able to add a separate layer of functionality which transforms TRAM to TRAM++. The processing of messages downstream and acks upstream in the no error scenario is shown in Figure 3 and for errors in Figure 4.

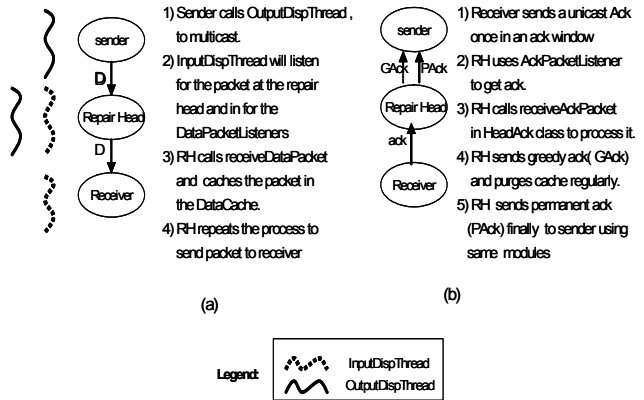
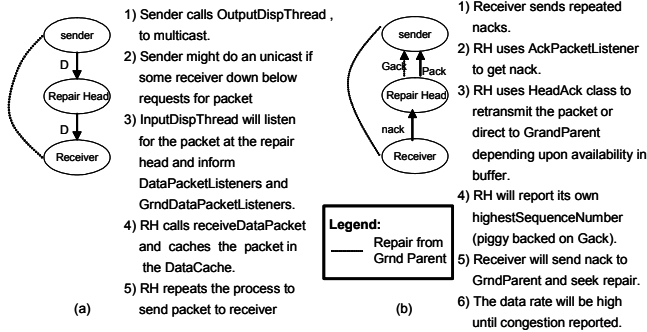


Figure 3. TRAM++ implementation for (a) upstream and (b) downstream message processing with no errors



**Figure 4. TRAM++ implementation for (a) upstream and (b) downstream message processing with message errors**

The ack processing mechanism of TRAM is modified for TRAM++. In TRAM++ as the ack (implicit nack) is received by the repair head, it checks which packets the receiver has asked for retransmission. If the packet is not in the cache, the repair head sends a packet called NOT\_REPAIRED to the receiver. This sub-message type forces the receiver to go for temporary re-affiliation. The data packet sent by the RH in response to the above type of nack, has a payload which only contains information about the higher-level repair head, also called *Grand Repair Head* (GRH). The receiver uses this information and sends a request for the missing packet to the new GRH. The GRH unicasts the requested data packet to the receiver. In TRAM, there is no notion of sending data through unicast. To enable this, we have added a message type called UCAST\_DATA. A flag is added to the ack packet of TRAM to differentiate between greedy and permanent ack types. The execution steps of TRAM++ with errors are shown in Figure 4. The downstream processing is identical to the TRAM case except that the sender may send a unicast if some receiver had directly requested a retransmission through the temporary re-affiliation process.

### 4.3. Experiments and Results

TRAM and TRAM++ are installed on a campus-wide WAN and experiments are conducted on the testbed. The experiments have two broad goals:

1. Evaluate the scalability of TRAM with the number of receivers and its robustness to different types of message errors.
2. Evaluate the improvement provided by TRAM++ in the event of failures and overhead incurred by TRAM++ under failure-free conditions.

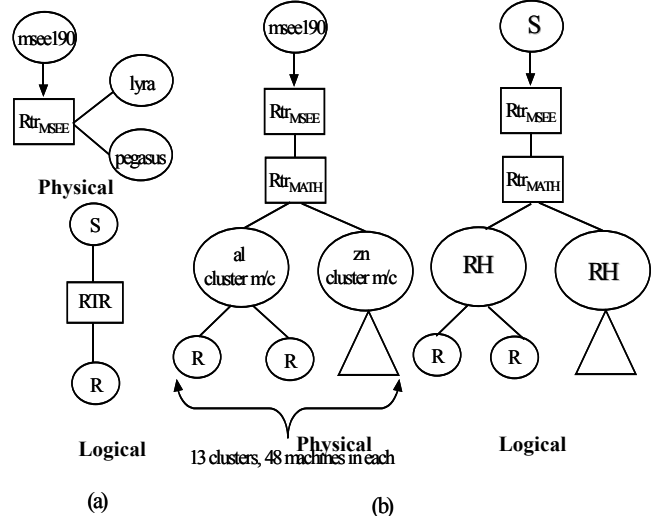
### 4.4. Testbed Setup

The testbed has a sender, multiple repair heads and varying number of receivers. The computing machines are

distributed across campus in the Mathematics (clusters *al-zn*), Electrical Engineering (named *pegasus* and *lyra*, together with its cluster) and Materials Science (*msee190*) buildings. The machine named *msee190* is always used exclusively as the sender. All the machines run RedHat Linux. It is important to note that the implementation is done and the experiments performed on a production campus-wide wide area network with normal traffic coexisting with the reliable multicast traffic.

**Table 1. Hardware Configuration**

MSEE machine (Sender)	Pentium IV, 1.5 GHz, 1 GB memory
EE machines	Pentium IV, 2.26 GHz, 1 GB, 533 FSB, 512 KB cache
Math machines	Pentium II, 450 MHz, 256 MB, divided into 13 clusters with 48 machines in each
Rtr <sub>MSEE</sub>	Cisco 5505 switch
Rtr <sub>MATH</sub>	Cisco 6509 switch
Links	Intra-cluster links 100 Mbps, inter-cluster links 1 Gbps



**Figure 5. Sample physical and logical configuration for deployment with (a) 2 hops (b) 4 hops between sender & receiver**

Different experiments are conducted with the receivers at different hops distance from the sender. In Figure 5, we show the physical and logical views of two different hop topologies. The physical topologies show the routers as well as the protocol participants (sender, receiver or RH), while the logical topologies only show the participants. A *hop* is defined as a link in the physical topology. This is different from the traditional definition of TTL in networking. For example, traversal from a receiver within a Math cluster machine to its RH in the same cluster does not decrement the TTL, but is considered a hop. We felt this to be reasonable because there is a cost of processing both at protocol participants and at the routers.

We use the reliable multicast infrastructure to send a high bandwidth Mpeg-2 video data feed from the sender to

the receivers. A 40 kbps feed is sent in 1316 bit payload packets, leaving space for the TRAM and the IP headers to fit within 1500 Ethernet MTU. The number of packets sent is at least 8000, with a larger number if the initial transients, which are discarded from our results, are longer. The parameters of the protocol set at the participants are shown in Table 2.

**Table 2. TRAM and TRAM++ configuration settings**

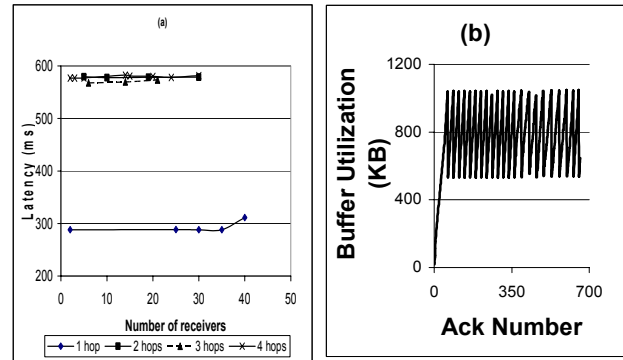
Parameter	Value	Parameter	Value
(Max, Min) data rate	(40 kbps, 1 kbps)	<b>TRAM:</b> Stable storage size at RH (TRAM)	1200 packets
<b>TRAM:</b> Low water mark: High water mark (# packets)	400:800	<b>TRAM++:</b> Stable storage size at RH	200 packets
<b>TRAM++:</b> Low water mark: High water mark (# packets)	32:120	<b>TRAM++:</b> Pruning of receiver (% of packets asking for reaffiliation)	0.08%

**Normal and Error Injection Runs:** A single run of the experiment is defined as the transfer of at least 8,000 packets of the video feed. A normal run is one where no errors are injected, though the variability of the environment may create congestion and transient instability such as spikes in latency. In the error injection runs, three kinds of message errors are simulated in the network – drops, delays and reordering. In one error injection run, a single receiver is identified as the faulty (or, malicious) one and the error injector works on its link. For message drops, a message is dropped on the downward link between the RH and a single receiver. For simulating message delays, the ack packet from the faulty receiver is delayed on the upstream link to the RH. For injecting message reorderings, we vary the inter-message gap ( $M_d$ ) between the messages which are to be reordered. Recollect that the ack window ( $W_a$ ) is the number of packets for which one ack gets sent. If  $M_d < W_a$ , then TRAM buffers the out-of-order packet and delivers it when an ordered sequence of packets can be created. This results in behavior identical to the error free case. If  $M_d \geq W_a$ , then the receiver sends a nack in the next ack window for the message whose place was taken by the out-of-order packet. This is identical to the message drop case and hence we do not separately show results for the message reordering case.

## 5. Evaluation of TRAM

### 5.1. Error-Free Case

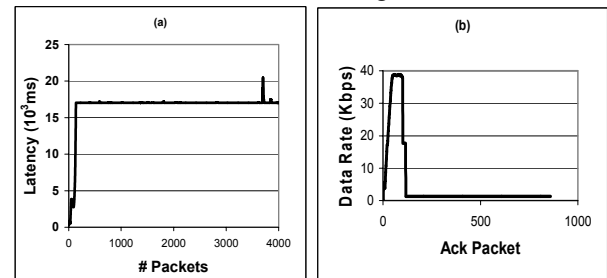
In this experiment, we investigate how TRAM scales with the number of receivers. The latency curves in Figure 6(a) show that the latency is substantially less for the single hop case since the message is on a direct connection and does not have to traverse a router. The latency for 2, 3 or 4 hops is comparable. None of the latencies become substantially worse with increasing number of receivers. The next experiment investigates the buffer utilization at the RH. TRAM is set up to use one-thirds of maximum buffer capacity (1200 packets) as the low water mark and two-thirds as the high water mark. Therefore, the buffer utilization oscillates between 400 and 800 packets, i.e., between 526.4 KB and 1.0528 MB. This is found to be true as the number of receivers is increased and is also independent of whether the utilization is at the RH or the sender. When we measure the data rate against time (or equivalently the number of the ack packet), it is found that, after the initial transient, the data rate picks up and tends towards the max data rate specified in the configuration.

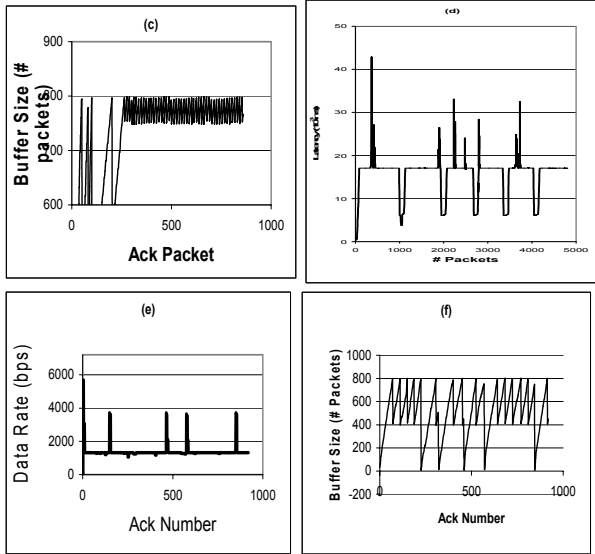


**Figure 6. Evaluation of TRAM in error-free case**

### 5.1. Error Injection Case

For TRAM, the effect of the error on the faulty receiver and the normal receiver will be identical, and therefore no distinction is made in the following discussion.



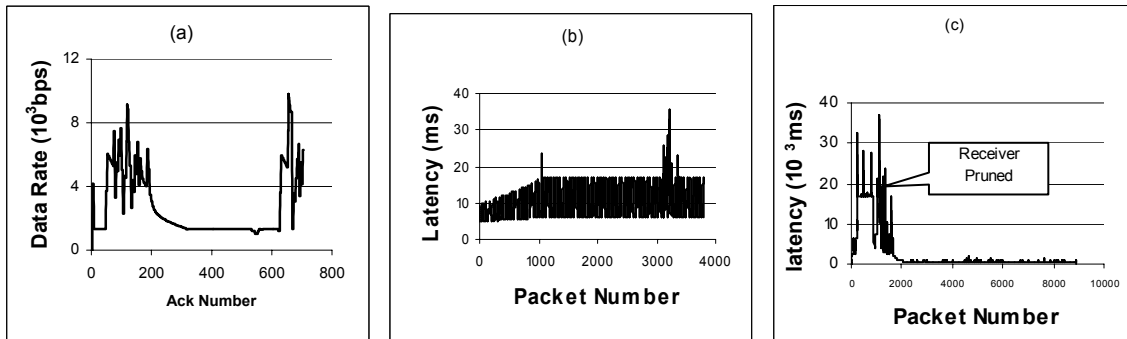


**Figure 7. Effect of message drop on TRAM: (a)&(d) Latency for 3 and 5 packets dropped out of 50, (b)&(e) Data Rates, (c)&(f) Buffer Occupancy**

Figure 7 shows the variation of latency, data rates, and buffer utilization with the number of receivers for two different message drop rates – 3 every 50 and 5 every 50 packets. The buffer utilization at the RH (Figure 7(c),(f)) shows an interesting trend. The purging of the buffer begins when its size reaches the high water mark (800 packets), and while under error-free conditions, the

purging would have been able to reduce the buffer utilization to 400 packets, here the reduction is only down to about 750 packets. If the drop rate is increased to 5 out of every 50 packets, the affected receiver is pruned. On pruning, the buffer is purged completely since the packets were being buffered to accommodate the receiver that just got pruned. Therefore the utilization comes down to zero. Immediately after pruning, the sender tries to increase the data rate and the utilization goes back to the usual oscillation between 400 and 800 packets till the pruning happens next.

In Figure 8, we show the effect of introducing message delays in TRAM. In the first experiment, a delay of 8,000 ms is introduced. The latency shows a regular spike of 8,000 ms and this spike repeats roughly every 32 packets since the ack window is set to 32 and a delay is introduced for every ack. The data rate shows congestion control at work. The sender tries to increase the data rate to the max rate, but is forced back because of the delayed ack. It reduces the data rate, but is able to sustain a rate greater than 1 kbps, and as a result, the receiver does not get pruned. The buffer utilization varies between the low and high water mark as in the error free case (Figure 6(b)). It is found that a delay of 10,000 ms causes pruning of the slow receiver. Once pruning of misbehaving receivers happens, the remaining receivers continue to see performance as in the error free case.



**Figure 8. Effect of message delay on TRAM: (a)-(b) Delay of 8000 ms, no pruning; (c) Delay of 10000 ms, pruning.**

## 6. Evaluation of TRAM++

### 6.1. Error-Free Case

The results of the scalability test of TRAM++ are shown in Figure 9. It is observed that the protocol is scalable like TRAM in the range under consideration (5-30 receivers). The variation in latency in this range is about 1.01%. The average latency in the range is 584.44 ms, which gives a 3.2% overhead over TRAM. As the

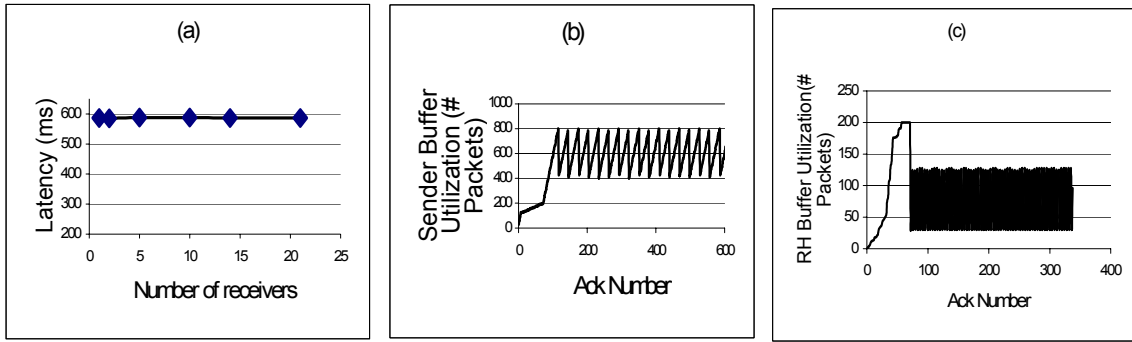
buffer constraint is varied in TRAM++, the relative overhead is found to remain constant in the error free case. This is expected since the buffer is not completely utilized in the absence of errors. The overhead is ascribed to three main reasons: extra message schema matching due to the introduction of new message types, aggressive cache pruning to satisfy the storage constraints at the RH, and additional control messages – two kinds of acks being sent upstream by the RHs.

Regarding the comparative resource usage, the CPU utilization is very close to zero for the maximum data rate



that the network infrastructure can support and therefore cannot form a meaningful point of comparison between the two protocols. The main memory utilization for both TRAM and TRAM++ vary between 7.5% and 30.0%. The sender buffer utilization curve shown in Figure 9(b) oscillates between 400 and 800 packets as in TRAM’s buffer utilization. The RH buffer utilization varies between the maximum pruning level (32 kB) and the high

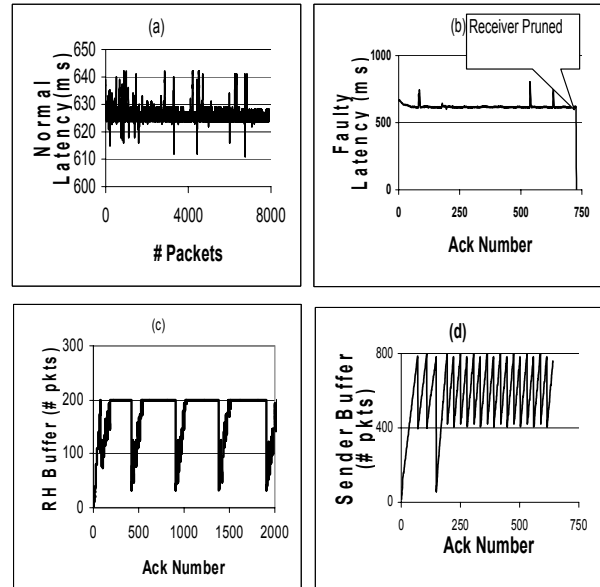
water mark (120 kB) and is shown in Figure 9(c). In the no error case, this buffer utilization does not depend on the number of receivers. A system designer can set the buffer upper bound knowing the stable storage constraints and TRAM++ will operate under the bound. The data rate supported by TRAM++ also approaches the max data rate parameter set in the system (40 kbps).



**Figure 9. Evaluation of TRAM++ in error-free case: (a) Receiver latency (b) Sender buffer utilization and (c) Repair Head buffer utilization for 27 receivers**

### 6.2. Error Injection Case

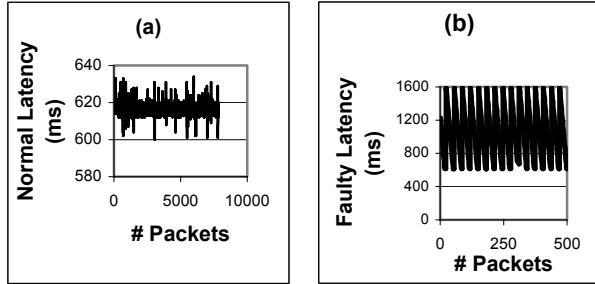
For message drops, we drop 5 packets out of every 50. The results are shown in Figure 10. It is observed that the malicious receiver gets pruned at around 650 packets, after which the system behaves as in the error free case. The sender’s buffer utilization once drops to zero because of purging of entire cache at the time of pruning. Then the utilization oscillates as usual between 400 and 800 packets, and the data rate also tends towards the maximum. However, the buffer utilization at the RH goes up to the maximum buffer space, before purging occurs. For 2 out of 50 packets being dropped, the receiver is repeatedly pruned and rejoins the multicast group. The normal or non-faulty node is not affected at all and its latency remains around the no error scenario value. This achieves the important design goal of TRAM++ of isolating the effect of a malfunctioning receiver to its part of the repair tree. Contrast this to the behavior in TRAM shown in Figure 7 where the latency of the normal receiver is shown to go above 10,000 ms for similar drop rates. We had the malicious receiver repeatedly reconnect to the multicast group to test the robustness of TRAM++ to this kind of malicious behavior.



**Figure 10. Evaluation of TRAM++ under message drop rate of 5 out of 50 packets.**

In the case of message delays, the pruning is found to happen for delays of 8000 ms and above. Figure 11 shows a scenario where pruning is not done (delay = 1000 ms). For the no-pruning scenario, the sender data rate and buffer utilization behave as in the error free case (see Figure 6(b) & Figure 9(b) respectively). The latency of the normal receiver remains unaffected though the sender detected congestion causes its latency to rise towards the end of the experimental run. The malicious receiver has a

saw-tooth latency pattern with the peak separated from the base by the delay amount.



**Figure 11. Effect of message delay (1000 ms) on TRAM++ (a) Normal receiver (b) Faulty receiver**

### 6.3. Highlights of Result

1. Both TRAM and TRAM++ scale well with respect to latency as the number of receivers is increased, up to a total of 30. Under error-free conditions, both were able to maintain a streaming video of data rate 40 Kbps, but not the next higher step of 60 Kbps.
2. TRAM is not successful in isolating the normal receivers from the effect of faulty or malicious ones. TRAM++ is able to achieve this through its protocol of differentiated acks and buffer management.
3. TRAM++ under a constraint of 16% of the TRAM buffer availability at the RH is able to maintain the end-to-end latency within 3.2% of TRAM for the error free case. TRAM++ achieves this without any additional memory overhead. In cases with message errors, the latency of normal receivers in TRAM++ is better by a factor of up to 30. TRAM++ is also able to prune malicious receivers faster because of local decision-making ability at the RH.

## 6. Conclusions

In this paper, we have presented an evaluation of a tree-based reliable multicast protocol called TRAM in error-free and message error prone environments. Two design improvements to TRAM are presented and the new protocol called TRAM++ is evaluated. The comparison brings out that both protocols are scalable in the range of receivers studied. If buffer utilization constraints are enforced on the intermediate nodes, then TRAM++ can operate under the constraints without substantial latency degradation compared to TRAM. It can localize the effect of a few slow or malicious receivers on the rest of the system. We are currently working on a distributed monitor architecture to detect failures or intrusions and work collaboratively with the TRAM++ participants to isolate the malfunctioning entity. The goal is to perform the diagnosis scalably and with low detection latency.

## Acknowledgements

This work is supported in part through a grant from the Purdue Research Foundation. Installation support for TRAM and TRAM++ was provided by Dale Talcott and Casey Carlson of Purdue ITAP.

## 7. References

- [1] S. Bhola, R. Strom, S. Bagchi, Y. Zhao and J. Auerbach, "Exactly-once Delivery in a Content-based Publish-Subscribe System". Dependable Systems and Networks (DSN'02), pp. 7-16, June 2002.
- [2] Dah Ming Chiu, Stephen Hurst, Miriam Kadansky and Joseph Wesley, "TRAM: A Tree-based Reliable Multicast Protocol", Sun Technical Report TR 98-66, July 1998.
- [3] Dah Ming Chiu, Miriam Kadansky, Joe Provino, Joseph Wesley and Haifeng Zhu, "Pruning algorithms for multicast flow control", Proceedings of NGC 2000 on Networked group communication, pp. 83-92.
- [4] Dah Ming Chiu, M. Kadansky, Joe Provino, J. Wesley, H. Bischof, Haifeng Zhu, "A congestion control algorithm for tree-based reliable multicast protocols", Proceedings of INFOCOMM 2002, pp.1209-1217.
- [5] H. Eriksson, "MBone: The Multicast Backbone", Communications of the ACM, pp.54-60, August 1994.
- [6] Mark Hayden, "The Ensemble System," Cornell University Technical Report, TR98-1662, January 1998.
- [7] T.Jiang, M.Ammar, and E.Zegura, "On the Use of Destination Set Grouping to Improve Inter-receiver Fairness for Multicast ABR Sessions", in Proceedings of INFOCOMM 2000, pp.42-51.
- [8] Sanjoy Paul and John C. Lin, "RMTP: A Reliable Multicast Transport Protocol", INFOCOMM 1996, pp.1414-1424.
- [9] Robbert van Renesse, Kenneth P. Birman, and Silvano Maffei, "Horus, a flexible Group Communication System," Communications of the ACM, April 1996.
- [10] S. Floyd, V. Jacobson, C.-G. Liu, S. McCanne, and L. Zhang, "A reliable multicast framework for lightweight session and application layer framing," IEEE/ACM Trans. Networking, vol. 5, Volume 5, Number 6, pp. 784-803.
- [11] Java Reliable Multicast Service. At: <http://www.experimentalstuff.com/Technologies/JRMS/>
- [12] R. Yavatkar, J. Griffioen, and M. Sudan, "A Reliable Dissemination Protocol for Interactive Collaborative Applications", In Proceedings of the ACM Multimedia '95 Conference, November 1995.
- [13] S. Floyd, V. Jacobson, C. Liu, S. McCanne, L. Zhang, "A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing", IEEE/ACM Transactions on Networking, December 1997, pp. 784-803.