

DyRAM : an Active Reliable Multicast Framework for Data Distribution

Moufida Maimour (moufida.maimour@ens-lyon.fr) and Cong-Duc Pham
(congduc.pham@ens-lyon.fr)
LIP, ENS de Lyon. France.

1997/01/23

Abstract. Group communications (multicast) are foreseen to be one of the most critical yet challenging technologies to meet the exponentially growing demands for data distribution in a large variety of applications of the Internet (such as grid computing, web applications and distributed simulations). When reliability is required, there is no straightforward solutions and meeting the objectives of reliable multicast is not an easy task. Active networks open a new perspective in providing more efficient solutions for the problem of reliability. In this context, routers are able to perform customized computations on the packets flowing through them. In this report, we propose a receiver-based (replier) local recovery multicast protocol with dynamic repliers elected on a per-packet basis. Designed to provide an efficient reliable multicast service without any cache facilities inside the network, our approach, uses low-overhead active services in routers. The current report addresses the design, evaluation and the implementation of an efficient and scalable reliable multicast protocol noted DyRAM standin for Dynamic Replier Active reliable Multicast.

Keywords: Active Networks, Reliable Multicast, Analysis, Simulation, Implementation

1. Introduction

Motivations behind multicast are to handle one-to-many communications in a wide-area network with the lowest network and end-system overheads, mainly by avoiding redundant traffic on physical links with packet duplications as close as possible to the final destination. Multicast is foreseen to be one of the most critical yet challenging technologies to meet the exponentially growing demands for data distribution in a large variety of applications of the Internet (such as grid computing, video-conferencing, web applications and distributed simulations). At the network level, IP multicast provides an efficient one-to-many IP packets delivery but without any reliability guarantees. However data dissemination applications such as distributed computing or interactive simulations usually require a reliable transfer that requires that all packets are safely delivered to the destinations.

1.1. ADDING RELIABILITY TO IP MULTICAST

Meeting the objectives of reliable multicast is not an easy task and, following the “IP vision”, IP multicast (RFC 1122 and [3]) has no mechanisms for reliability. Therefore reliability has to be added at a higher layer. The problem of reliability in multicast protocols has been quite widely covered during the last 10 years. Early reliable multicast protocols use an end-to-end solution to perform the loss recovery. Most of them fall into one of the following classes: sender-initiated, receiver-initiated and receiver-initiated with local recovery protocols. In sender-initiated protocols, the sender is responsible for both the loss detection and the recovery. These protocols usually do not scale well to a large number of receivers due to the ACK implosion problem. Receiver-initiated protocols move the loss detection responsibility to the receivers. They use NACKs instead of ACKs. However they still suffer from the NACK implosion problem when a large number of receivers have subscribed to the multicast session. In receiver-initiated protocols with local recovery, the retransmission of a lost packet can be performed by some other nodes in the multicast tree [5, 14, 18, 13, 15, 7]. Most solutions now propose local recovery to avoid the implosion of control messages at the source and to reduce the recovery latency. There are basically two possibilities for enabling local recoveries: (i) use some receivers or dedicated servers as repliers (in some cases with a router-assisted solution), or (ii) use network elements such as routers for caching. In the first category, the replier could be any receiver in the neighborhood (SRM [5]), a designated receiver (RMTP [14], TMTP [18], LMS [13], PGM [15]) or a logging server (LBRM [7]) in a hierarchical structure.



1.2. TOWARDS ACTIVE RELIABLE MULTICAST

Local recovery appears to be the most promising choice for achieving reliability on an open Internet network. Using a replier has the main advantage of requiring a memory usage as low as possible within network elements and is potentially more scalable. The choice of the replier can be done in several ways. Protocols that use some kind of router assistance are for instance LMS [13] and PGM [15]. LMS consists in adding more topology information with little help from routers. With some specific forwarding functions within routers, the protocol elects a designated replier (leader) for each subtree to respond to the requests made by the end hosts. PGM also uses some router assistance (not yet active router but rather PGM router) to discover and elect repairers that must be on the path towards the source.

Going a step further, and gaining in generality and flexibility, the use of active network concepts [17] where routers themselves could contribute to enhance the network services by customized functionalities have been proposed in the multicast research community. Recently, the use of these active network concepts has been proposed in many research areas including multicast protocols [9, 8, 2], distributed interactive simulations [4] and network management [16]. Contributing mainly on feedback implosion problems, retransmission scoping and cache of data, active reliable multicast offer a general and flexible framework for customized functionalities in network protocols. ARM [9], AER [8] and RMANP [2] are protocols that use cache of data packets to permit local recoveries and advanced NACK suppression strategies. ARM (Active Reliable Multicast) and AER (Active Error Recovery) are two protocols that were recently proposed in the research community and that use active services within routers.

The approach we propose use local recoveries performed by repliers elected amongst a sub-set of receivers. Instead of an approximate solution based on timers as in SRM, or a complex DLR discovery as in PGM, the elected replier in our case is dynamically determined at each lost packet (without much overhead as it will be described later on), and not determined at the beginning of the multicast session, thus justifying the “dynamic replier” property of our active reliable multicast protocol (noted DyRAM for *Dynamic Replier Active reliable Multicast*). Since a specific replier can be chosen for each lost packet, it is therefore possible to have several logical subtrees at the same time for the recovery process. In addition, this very dynamic choice of the replier provides load balancing features that reduces the receiver overhead. As opposed to LMS and PGM, DyRAM uses intensively the concept of active networks with customized services within routers for implementing several advanced mechanisms including the replier election. This choice of active networking is motivated by the fact that the deployment on-the-fly of customized services appears to be more general and flexible than a solution based on a dedicated, static router (PGM routers for example). However, this choice leads to some design constraints where the active services must incur as low as possible execution/memory overhead within routers since the service will mainly be executed in software.

Therefore we can summarize the motivations behind DyRAM with the following design goals: (i) to minimize active routers load in order to make them supporting more sessions (mainly in unloading them from the function of data caching performed in ARM), (ii) to perform a more efficient replier election procedure with active services, (iii) to dynamically distribute the recovery task among the downstream receivers to do not overload one replier and (iv) to incur as low as possible execution/memory overheads.

The rest of the paper is organized as follows. Section 2 gives an overview of the DyRAM protocol and its active services are described. A performance evaluation of DyRAM using simulation is the object of section 3. In section 4, we propose our active architecture more suitable to reliable multicast protocols. Section 5 shows how DyRAM is implemented in an active execution environment and preliminary results will be presented in measuring the raw processing time required for enabling low-latency multicast communications on the Internet. Section 6 concludes.

2. DyRAM General Features

2.1. DESCRIPTION

DyRAM is a reliable multicast protocol with a recovery strategy based on a tree structure built on a per-packet basis with the assistance of routers. It is an incremental step from the existing propositions but it has been designed to provide low-overhead active functionalities with a dynamic replier election on a per-packet basis. The protocol uses a NACK-based scheme with receiver-based local recovery where receivers are responsible for both the loss detection and the retransmission of repair packets when it is possible. In order to perform flow and congestion control as well as memory management in addition to a more efficient replier election, ACKs are piggybacked in the NACKs. In the absence of NACKs, ACKs are also periodically piggybacked by special messages called “Congestion Reports” (CRs). These CRs contain the necessary information for the congestion and flow control. Routers play an active role in DyRAM which consists (in addition to the services related to congestion avoidance and flow control) in the following active services:

- Duplicate NACK suppression in order to limit the NACK implosion problem.
- CRs aggregation thus avoiding the CR implosion at intermediate nodes and at the source.
- Subcast of the repair packets only to the relevant set of receivers that have experienced a loss. This helps to limit the scope of the repairs to only the affected receivers.
- Dynamic replier election which consists in choosing a link as a replier to perform a local recovery from a downstream receiver instead of caching data packets at the routers. Dynamic election provides robustness to host and link failures.
- Early loss detection which is very helpful in providing a low recovery delay.

In DyRAM, the source sends data packets to the multicast address subscribed to by all the receivers. The core functionalities of IP multicast are assumed to deliver the packets to the receivers in a best effort manner. A receiver detects a loss by sequence gaps and upon the detection of a loss, a receiver immediately sends a NACK towards the source and sets a timer. Since NACKs and repairs may also be lost, a receiver will re-send a similar NACK when the requested repair has not been received within the timeout interval. Practically, the timeout must be set to at least the estimated round trip time (RTT) to the source. In order to detect the loss of the last packets (no gap in sequence number in this case), a receiver will request the next packet which is supposed to be received if no data has been received within a predefined receiving time window. This window width depends on the transmission rate of the source. The last data packet is easily identified as the EOF (end-of-file) packet which will be sent by the source as the last one. Upon the reception of a NACK packet, the source sends the repair packet to the multicast address.

In order to limit the processing overheads of duplicate NACKs and to avoid the corresponding retransmissions, the source, the active routers and the receivers (remember that DyRAM uses repliers among receivers) ignore similar NACKs for a certain period of time. According to this rule, a NACK will be said “valid” if it is received for the first time or no similar NACKs are received during this discard time window. How this discard window and the other timeouts are adjusted and how the different round trip times (RTTs) are estimated is addressed in section 2.5.

To avoid congestion, DyRAM adjusts its rate periodically based on the reception of CRs from the receivers which contain mainly information about RTT variation. To avoid the problem of rerouting paths, we adopt a hop-by-hop congestion control approach where communications between adjacent nodes including the routers, are performed. Active networks allow us to perform on a per link basis dialogue. A minimum and a maximum rates (r_{min} and r_{max}) are set depending on the application. Initially




Figure 1. Packets structure.

the source starts to send data packets with a rate equal to the minimum rate r_{min} . Then the source tries to increase its rate if no congestion indication is received without exceeding the maximum rate r_{max} . The source uses the information fed back in the CRs and NACKs to update its rate. For instance, the RTT variation field of every CR is used by the source to compute the queue size variation per packet during the previous period. For more details about the congestion control in DyRAM, the reader can refer to [12].

2.2. PACKETS STRUCTURE

All the used packet types contain in their header the multicast address, the source address and the active service identifier (*SVC*) to be performed on it (see Figure 1). A data packet is labelled uniquely by a sequence number and contains a dedicated field (*isR*) to distinguish an original transmission from a repair. A data packet contains also the current emission rate which is mainly useful to be known by the receivers in order to correctly update timeouts.

A NACK or a CR packet include in addition to their originator address, their last ordered (*lo*) and last received (*lr*) data packet. Moreover, a receiver will report its RTT to the source and to its active router using the *RTT2S* and *RTT2P* fields contained in NACK and CR packets. Whereas a NACK contains the sequence number of the requested data packet, A CR contains two other fields which are the CRs current period and an aggregation of the RTT variation experienced by the subtree from which the CR is received.

For the purpose of flow control, every receiver includes in the *MaxSeq* field of a CR or a NACK, the maximum data packet sequence number that corresponds to its available buffering means. The *MaxSeq* field in a CR or NACK packet is used by the source to regulate the transmission flow. An active router includes in the CR to be sent upstream to the source, the minimum *MaxSeq* value among those reported by the different CRs received from downstream. In this way, the source ends up by receiving the minimum *MaxSeq* reported by the different receivers. The source will send with its current transmission rate, data packets with sequence number not greater than this minimum *MaxSeq* value.

The *lo* field is also used to release buffers at both the sender and the receivers. With the proposed mechanism of aggregating the CRs, the source ends-up by receiving the minimum *lo* of the whole group. This is an acknowledgement that all data packets up to *lo* have been correctly received by all the receivers. The source could then release buffers of data packets up to the *lo*th one. The source includes the current *lo* of the whole group in the *MaxSeq* field of the data packets it sends. A receiver on the reception of a data packet with a given *MaxSeq* field, would update its flow control window accordingly. Since a receiver would keep data packets in order to be able to act as an eventual replier, this mechanism is helpful to releasing memory. In fact, a receiver could remove from its memory all the data packets with a sequence number less or equal to the specified *MaxSeq*.

2.3. DYRAM ACTIVE SERVICES

2.3.1. Router's Soft States

Within active routers, the different active services can be implemented simply by maintaining information about the data packets, NACKs and CRs received. This set of information is uniquely identified by the session source and the multicast address. For each received NACK, the router creates or simply updates an existing NACK state (NS) structure which has a limited life time. A NS is removed on the reception of the corresponding repair. Hence its life time can be estimated as the maximum expected time required by the corresponding repair to be received by the router. This structure maintains for every nacked packet, a subcast list (*subList*) that contains the numbers of links from which a NACK has been received.

DyRAM also enables the routers to keep track of the received data packets in order to quickly detect packet losses occurring from upstream and affecting all its subtree. This can be done simply by maintaining a track list (TL) structure for each multicast session handled by the router. A TL has three components:

- *lastOrdered* is the sequence number of the last data packet received in order.
- *lastReceived* is the sequence number of the last received data packet.
- *lostList* contains the list of data packets not received by the router with sequence number greater than *lastOrdered* and less than *lastReceived*.

In order to perform quicker, accurate and optimized replier elections (as will be seen in section 2.3.5), we keep within a router a link state structure (LS) which contains for every downstream link, the following¹:

- *lastOrdered* is the sequence number of the last data packet received in the order by this link (by all the receivers behind this link).
- *lastReceived* is the sequence number of the last received data packet by this link.
- *weight* is an assigned weight to this link which is proportional to its capacity to be a replier.

2.3.2. NACK suppression and CR aggregation

On receipt of a valid NACK for a data packet, a router creates the corresponding NS structure and initializes a timer noted NST (NACK Suppression Timer). During this timer, all subsequent NACK packets received are used to appropriately update the LS and NS structures and are dropped afterward. According to our definition of a valid NACK in section 2.1, a NACK is considered as valid if when it is received there is no active NST.

The LS structures are also updated on the reception of the CRs. All the CRs received from the downstream links are aggregated in one CR to be sent on the upstream link. The aggregated CR contains the sequence number of the minimum last ordered and the maximum last received data packets among those reported by the CRs received from downstream. For the RTT variations reported by the CRs from downstream in a multi-path connection, we use:

$$\Delta\tau = \Delta\tau_{up} + \text{Max}_i(\Delta\tau_i)$$

where $\Delta\tau_i$ are the reported RTT variations from the downstream links and $\Delta\tau_{up}$ is the RTT variation experienced by the upstream link.

¹ These information are retrieved from the CRs and NACKs received during the multicast session as will be explained in the next sections.

2.3.3. Subcast Functionality

The subcast list in the NS structure is an important component for the subcast functionality. This list contains the set of links (downstream or upstream) from which a NACK has been received. When a data packet arrives at an active router it will simply be forwarded on all the downstream links if it is an original transmission. If the data packet is a repair packet the router looks for a corresponding NS structure and would send the repair on all the links that appear in the subcast list.

2.3.4. Replier election

The “replier election” is a key functionality which is specific to DyRAM. It allows for local recovery with support from the active routers instead of loading them by heavy caching functions. On the receipt of a NACK for a data packet from a given link, a router would update the corresponding LS structure when necessary. If the received NACK is valid, an NST is initialized and an NS structure is created. Additionally, the router verifies (using the LS structure) if the requested data packet has already been acked by a downstream link. If so the NACK will be immediately forwarded downstream on one of the links that appear to receive this data packet. Otherwise a replier election phase is initialized by setting a timer noted DTD (Delay To Decide) in order to collect during this time window as much information as possible about the links affected by a loss (updates of the subcast list and the LS structure).

If the received NACK is not valid, the router updates the corresponding NS structure. If no DTD is active for the requested data packet then the NACK is simply dropped. Otherwise if all the downstream links appear to be in the subcast list then the NACK is immediately forwarded up. The LS structure could be useful to perform quicker and more accurate replier election. When a received CR (or a NACK) piggybacks an ACK for the lost data packet, the corresponding link is elected immediately as a replier and the DTD timer is canceled. If no event has triggered the replier election, this latter is performed when the DTD timer expires. A router would choose one of the downstream links which is not in the subcast list.

Once the election is completed, the router will forward the NACK downstream on the elected link. If this link ends to a receiver, then this latter will be the elected replier. Otherwise it ends to an active router which in turn will perform a replier election and so on until the NACK reaches a receiver. The elected replier would then unicast the repair packet to its upstream router which would in turn subcast the repair packet on all the links in the subcast list. To avoid the processing of the NACKs for a data packet for which a repair has already been sent, an active router ignores requests for this data packet during a given period of time.

2.3.5. Optimizing the replier election

The elected replier in our case is dynamically determined at each lost packet, and not determined at the beginning of the multicast session as in LMS [13]. The replier link in LMS is only updated when the group membership changes or a receiver is no longer willing to act as a replier. Thus justifying the “dynamic replier” property of our protocol. To avoid overloading a particular downstream link, the router always try to choose a different link from the previously elected one (if available) by respecting a ring order among them, thus realizing when possible a load balance (in addition to providing robustness to replier and link failure).

Using the *weight* field of the LS structure, a more sophisticated replier election could be performed. In fact by mean of the *weight* field, a “replying capacity” is assigned to each downstream link. When the choice between multiple downstream links is present, we try always to choose the link that appears to be the most “capable” one in a way so the number of times a link is elected will be proportional to its assigned weight. To estimate the weight to be associated to a given link, we use the following:

$$weight = \kappa (-RTTvar) + (1 - \kappa) \frac{1}{RTT \times (lr - lo + 1)} \quad (1)$$

where κ is a weighting parameter to be chosen in $]0, 1[$. $RTTvar$ varies inversely with the available bandwidth on the considered link. The RTT gives an idea on the potential replier distance to the active router. As for $(lr - lo + 1)$, it can be seen as an estimation of the loss rate experienced by the receiver located downstream this link.

2.3.6. Early Loss Detection

Earlier reliable multicast protocols put the burden of both the loss detection and the recovery at the sender side. These protocols are TCP-like and use ACKs combined with timers to detect losses. These ACK-based schemes violate the IP multicast Model [3] stating that the source is not aware of the identity of the receivers. The use of NACKs instead of ACKs moves the loss detection responsibility to the receivers. A receiver detects a loss on receipt of an out-of-order data packet or on a timeout expiration. When a loss is detected, the receiver sends a NACK to the source requesting the lost data packet. In DyRAM, like the receivers, active routers are also responsible for the loss detection. In fact, active routers are capable to detect packet losses and consequently generate corresponding NACKs to be sent to the source.

An active router in a similar way to a receiver, would detect a loss when a gap occurs in the data packet sequence or on timeout expiration. The data service, in the absence of losses, simply consists in updating the track list. When a loss occurs, the router immediately generates a NACK packet towards the source. If the router has already sent a similar NACK for a lost packet then it will refrain from sending a NACK for a given amount of time (based on a timer). During this period, all NACK packets received for this data packet from the downstream links are ignored.

2.4. DYRAM MAIN ALGORITHMS

This section presents in a synthetic manner the algorithms performed by the active routers in the multi-cast tree on a data and a NACK packet. RTT estimation, flow control and congestion avoidance related algorithms are not considered here.

2.4.1. The Data Packet Service

When an active router receives a data packet (DP), it will execute the data packet service described below:

```

update the trackList of the session (TL)
if DP is a repair then
    look for the corresponding NACK state structure NS
    send DP through links in NS.subList
    free NS
else
    forward DP on all the downstream links
    for each  $DP_{lost}$  with  $((DP_{lost}.seq > TL.lastOrdered + 1)$  and  $(DP_{lost}.seq \in TL.lostList))$  do
        send a NACK( $DP_{lost}.seq$ ) to the source
        create or update the corresponding NS structure  $NS_{lost}$ :
         $NS_{lost}.seq \leftarrow DP_{lost}.seq$ 
         $NS_{lost}.subList \leftarrow$  all the links downstream
    end for
end if

```

2.4.2. The NACK Packet Service.

Provided that $B \geq 2$ where B is the number of the router's downstream links (if $B = 1$, then the NACK is simply forwarded upstream), an active router receiving a NACK packet from the link number i will perform the following actions:

```

update the corresponding LS structure if necessary:

```

```

 $LS_j.lastOrdered \leftarrow NACK.lo$ 
 $LS_j.lastReceived \leftarrow NACK.lr$ 
if NST is active then
  update the corresponding  $NS$  structure:
   $NS.seq \leftarrow NACK.seq$ 
   $NS.subList \leftarrow NS.subList \cup \{i\}$ 
  if DTD timer is active then
    if all the downstream links are in  $NS.subList$  then
      forward the NACK upstream toward the source
      cancel the corresponding DTD timer
    else if  $\exists j \neq i, LS_j.lastOrdered \geq NACK.seq$  then
      cancel the corresponding DTD timer
      forward the NACK on link  $j$  downstream
    end if
  end if
  drop this NACK
else
  initialize NST
  create the corresponding  $NS$ :
   $NS.seq \leftarrow NACK.seq$ 
   $NS.subList \leftarrow NS.subList \cup \{i\}$ 
  if  $(\exists j \neq i, LS_j.lastOrdered \geq NACK.seq)$  and (DTD is active) then
    cancel the corresponding DTD timer
    forward the NACK on link  $j$  downstream
  else
    initialize a DTD timer
  end if
end if

```

2.5. ESTIMATION OF THE RTTs AND TIMEOUTS SETTING

Many of the DyRAM's functionalities depend on appropriately setting the timers. These timers values depend on the different RTTs between the tree nodes. In this section, we first report our method to estimate the RTTs and secondly present how the different timeouts are set.

2.5.1. RTT estimation

A straightforward solution to estimate the RTTs between each receiver and the source consists in sending ping messages periodically from the receivers to the source. This latter will be overwhelmed in the presence of a large number of receivers. In DyRAM instead of estimating directly the RTTs between the source and each receiver, we begin by estimating the RTTs between every node and its parent in the multicast tree. To illustrate our method we show in figure 2, one path of a multicast tree that connects the source (labeled by 0) to one of the group receivers (with the label n). Intermediate nodes which are labeled by numbers from 1 to $n - 1$ are the routers. The generalization of our method to the other receivers and intermediate nodes is straightforward.

Heartbeat messages (HB and HB_RESP) are periodically exchanged between each node i and its parent ($i - 1$) in the same way as the ping/pong messages, instead of sending the HB messages until the source. The reception of a HB by a router triggers the emission of a HB message to its parent in the multicast tree in addition to responding with a HB_RESP message.

For the following, we note $\theta_{i,j}$ the RTT between the i th and the j th node. Recall that $\forall i, j : \theta_{i,j} = \theta_{j,i}$. Consequently, $\theta_{i-1,i}$ is the RTT between the node i and its parent and $\theta_{0,i}$ is the RTT between a node i and the source.

A node i , in order to compute its RTT to the source, first computes its RTT to its parent node by sending an HB message timestamped with the emission time. The parent node, On the reception of the




Figure 2. RTT estimation

HB message, responds with an HB_RESP that contains the parent's i th node RTT to the source already computed. The node i , on receipt of the HB_RESP message is able to compute its RTT to the source as the sum of the RTT of its parent to the source and its own RTT to its parent as follows:

$$\forall i > 0 : \theta_{0,i} = \theta_{0,i-1} + \theta_{i-1,i} \quad (2)$$

Initially and for a first estimation, the source multicasts a special message HB_RESP_S which contains its own RTT to itself (of course we have $\theta_{0,0} = 0$). The node number 1 will update its $\theta_{0,1}$ using (2) and forwards it downstream. The node number 2 will in turn, update its $\theta_{0,2}$ using the same equation and forwards it downstream and so on (figure 2). Afterwards the use of the HB and HB_RESP messages will be sufficient to estimate the different RTTs.

2.5.2. Timeouts Setting

If all routers at the edge of the core network are active, our protocol performs well with respect to different metrics as will be seen. The feedback implosion is avoided by performing NACK suppression at the routers. Duplicate repairs disappear as a consequence of the NACK suppression service. In fact, if all the routers are active and have sufficient memory to create an NS structure for each received NACK without replacing anyone of them unless the corresponding repair is received, then the source or an elected replier will receive at most one NACK per loss. Otherwise, the sender or a receiver respond to the first NACK, then subsequent NACKs for this packet are ignored for a fixed amount of time. This discard time window must be well adjusted. A very small time window will produce duplicate repairs and a very large one will increase the loss recovery delay due to the fact that lost repairs are not quickly detected. This time window depends on the farthest receiver in the group. For our simulations this time window is set to $(\max(RTT_{down}) - \min(RTT_{down}))/2$ where $\max(RTT_{down})$ and $\min(RTT_{down})$ are the maximum (respectively minimum) measured RTT between the source and any other receiver in the multicast tree.

We divide by 2 to have a mean value². This timeout as has been discussed is not really required because of the NACK suppression, hence it can be set to a minimum value.

Since NACKs and repairs may also be lost, a receiver will re-send a similar NACK when the requested repair has not been received within a timeout interval. This timer used by a receiver to retransmit a new NACK when a repair is not received is at least set to its RTT to the source. To detect the loss of the last data packets, a receiver uses a timer which is set to at least $(1/rate_{current})$ where $rate_{current}$ is the current transmission rate of the source. The current rate is known by each receiver thanks to a dedicated field in the data packet header (section 2.2). This field is also used by the congestion avoidance algorithm associated to DyRAM [12].

As described in section 2.3.4, active routers perform the election of a replier link on a per-packet basis. In order to take the best decision about which links may be the replier link, the DTD timer has to be well chosen. Neither a very small nor a very large one is recommended. With a very small one, a router will not be able to elect the appropriate link as a replier. A very large DTD will increase without justification the recovery latency. A good value for the DTD timer depends on how soon the router expects to receive all the NACKs from the downstream links. The DTD timer is set using :

$$DTD = \delta DTD_{min} + (1 - \delta) DTD_{max}$$

where

$$DTD_{min} = (\max(RTT_{down}) - \min(RTT_{down}))/2$$

$$DTD_{max} = DTD_{min} + period$$

where DTD_{min} (DTD_{max}) is the minimum (maximum) estimation of the required time to receive the necessary information to be able to decide which link will be the elected one. *period* is the current period for sending CRs [12]. $\max(RTT_{down})$ ($\min(RTT_{down})$) is the maximum (minimum) RTT between this active router and any receiver in its subtree.

An active router *A* ignores requests for a data packet if a repair has recently been forwarded for approximately $RTT(A, R_x)/2$ where R_x is the furthest receiver in the subtree of *A*. In addition, if the router is loss detection-capable, then a NACK could be resent if a requested repair is not received yet. When a NACK is sent, a router sets a timer whose the expiration triggers the emission of a new NACK. This timer is set (as for the receivers) to at least its RTT to the source. During this time, all NACK packets for this data packet from the downstream links are ignored.

3. Simulation Results

We implemented a simulation model of DyRAM (in the PARSEC language developed at UCLA [1]) and evaluated its performances on a network model. The network model considers one source that multicasts data packets to *R* receivers through a packet network composed of a fast core network and several slower edge access networks. We will call *source link* the set of point-to-point links and traditional routers that connects the source to the core network. Similarly, a *tail link* is composed of point-to-point links and routers connecting a receiver to the core network (see Figure 3). We consider that the path to every receiver contains l_b backbone links. We only consider active routers at the edge of the core network. This is due to the fact that the core network is reliable and a very high-speed network. Adding complex processing functions inside the core network will certainly slow down the packet forwarding functions.

At the beginning, simulations are performed for a basic DyRAM protocol without the loss detection service which is afterwards integrated to perform an other set of simulations. In both of the two sets of simulations, the replier election is performed with respect to a ring order without any optimization. it is worth mentioning that the congestion avoidance is turned off for these simulations. The purpose is to be

² An EWMA filter can be used.

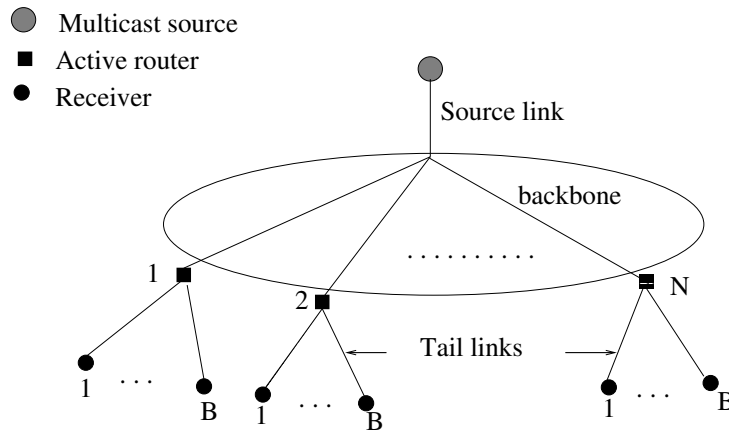


Figure 3. Network Model.

able to experience losses. Simulation results related to the congestion avoidance algorithm can be found in [12].

3.1. METRICS

To evaluate our protocol, we have defined a set of metrics. Firstly, M_S is defined as the number of retransmissions from the source per packet and gives an idea on the load at the source.

$$M_S = \frac{\text{Number of retransmissions}}{\text{Number of sent packets}}$$

To quantify the load at the network level, we use BW which is the average bandwidth consumed per link during the multicast session:

$$BW = \frac{BW_{NACK} + BW_{Data}}{N \cdot B + l_b \cdot (N + 1)}$$

where BW_{NACK} and BW_{Data} are respectively the bandwidth consumed by the NACK packets and the data packets during the multicast session: $BW_{NACK} = \sum_{i=1}^{NackNb} l_{NACK_i}$ and $BW_{Data} = \sum_{i=1}^{DataNb} l_{DP_i}$ where l_{NACK_i} and l_{DP_i} are respectively the number of links crossed by the i th NACK packet and the i th data packet. The third metric is the completion time per packet which is the required time to successfully receive the packet by all the receivers (can be seen as the latency).

For all the simulations, A NACK and a data packet are considered to be of 32 and 1024 bytes respectively. All simulation model values are normalized to the NACK transmission time. For the processing overheads at the routers, we assume that both NACK and data packets are processed in 32 time units. These values are derived from measures in [9]. In our simulation models, we have taken into consideration the fact that the repairs may also be lost.

In the simulation results, first we will show the benefit of performing the local recovery from the receivers. Then a comparison between our approach (no cache in routers) and ARM (with cache in routers without repliers) will be presented. Finally, we will consider the case when DyRAM also benefits from some cache amount in the routers. Due to space limitation, some curves are not presented, for more curves the reader can refer to [11].

3.2. SIMULATION RESULTS WITHOUT THE LOSS DETECTION SERVICE

The network model used in this first set of simulations considers active routers only at the receivers side. Each active router A_i is responsible of B receivers noted R_{i1}, \dots, R_{iB} forming a local group. In our loss model, we consider both the spatial and the temporal correlation of data packet losses. The spatial

correlation is introduced by considering a per-link loss rate. The core network is considered reliable as mentioned previously. For the other links (the source link or the tail links), the loss probability is noted p_l . Therefore, the end-to-end probability of a packet loss perceived by a receiver is $p = 1 - (1 - p_l)^2$. The losses at the tail links are assumed to be mutually independent.

3.2.1. Local Recovery from the Receivers.

Figure 4 plots for DyRAM the number of retransmissions (M_S) from the source as a function of the number of receivers, with and without local recovery, for 5% and 25% loss rates. These results have been obtained from simulations of a multicast session with 48 receivers distributed among 12 local groups. The curves show that the local recovery introduces less load at the source as the number of receivers increases, especially for high loss rates. Even not shown, the consumed bandwidth and the completion time are reduced in about the same proportions.




Figure 4. Load at the source in DyRAM with and without local recovery (a) $p = 0.05$ and (b) $p = 0.25$.

Putting the recovery process in the receivers requires at least 2 receivers per active router otherwise local recovery can not be realized. Therefore the local group size (B) is an important parameter. In order to study the impact of B , simulations are performed with 48 receivers distributed among groups of different sizes. Figure 5 compares the per-link consumed bandwidth ratio of DyRAM to an other approach without any local recovery facilities as the group size is varied. As the receivers number per group increases, the consumed bandwidth is reduced with larger ratios for large loss rates. In fact when the group size is larger, there are more chances that the members of the group can recover from each other. For instance, Figure 5 shows that with only 6 receivers per group and local recovery we achieve a gain of 80%.

3.2.2. DyRAM vs. ARM.

In this section, we compare our approach to ARM. Figure 6 shows for different loss rates, the load at the source in term of number of retransmissions for ARM and DyRAM as a function of the percentage of the available cache at the ARM routers. Without any cache at the routers, DyRAM always performs better than ARM. When ARM benefits from caching at the routers, it performs better only for very low loss rates or large cache sizes at the routers. However for large loss rates (e.g. 25%), ARM requires more than 50% of cache at the routers to perform better than DyRAM. Our approach is more promising because even without any caching overhead at the routers DyRAM still has better performances than ARM when the latter can use more than 50% of cache.

3.2.3. DyRAM combined with cache at the routers.

The results we present now assume that some cache memory is available in routers in addition to the DyRAM native local recovery strategy. Figure 7 plots for different loss rates the gain in consumed




Figure 5. Bandwidth gain ratio in DyRAM as a function of the group size.




Figure 6. DyRAM vs. ARM : load at the source.

bandwidth achieved by ARM and DyRAM when the cache size is varied over a simple ARM without caching in routers. The main results are that with cache facilities at the routers for both protocols, ARM never performs better than DyRAM even for low loss rates. When no cache is available at the routers, DyRAM outperforms ARM and this for any loss rate. As p increases, DyRAM performs better and better. In fact for 25% loss rate, ARM needs more than 50% of cache to achieve the same level of performances than DyRAM without any caching at the routers. To obtain a gain of 20%, ARM must benefit from 40% of cache. Instead, DyRAM without any cache already performs 40% better and with 40% of cache, it performs 60% better. One different way to look at the figures is that if DyRAM has only 40% of cache available then it already performs better than ARM with more than 60% of cache (for 5% loss rate). With 80% of cache, DyRAM has the same level of performances than ARM with 100% of cache available. Stated differently, about 80% of cache at the routers gives DyRAM the same features than a protocol with an infinite storage capacity.

3.3. ADDING THE LOSS DETECTION SERVICE TO DYRAM

Here we will consider the loss detection service in our simulations. Although simple, moving the loss detection into routers arises many questions such as “where to place such a service so it will be more efficient ?” and “what is the overhead of such an additional service?”. A previous analytical study [10] of this active service has shown that one must be careful when doing so in order to get any real benefit. The analytical study was performed using a simple network model with a two level multicast tree and consists in a delay analysis.

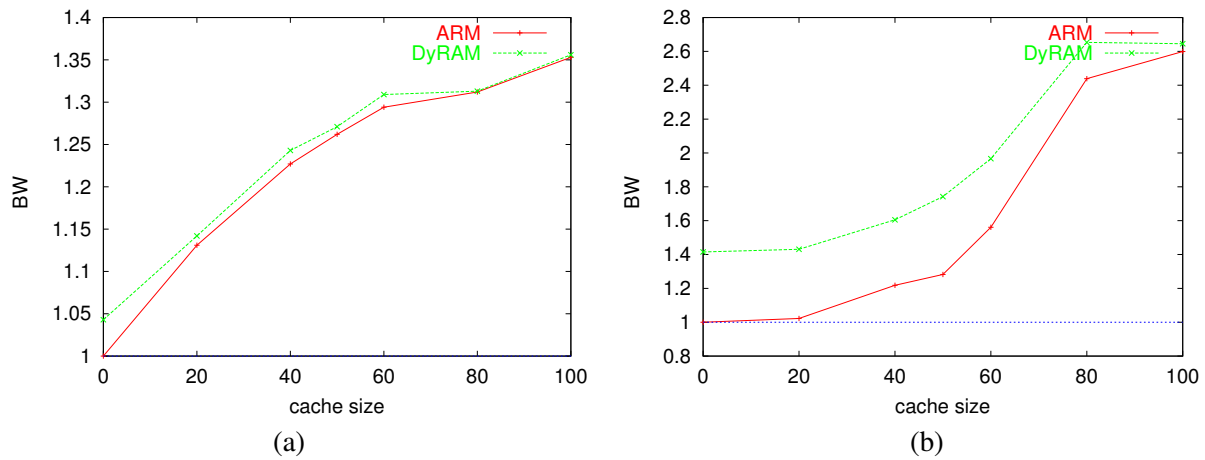


Figure 7. Gain obtained by DyRAM and ARM when both benefit from cache in the routers (a) $p = 0.05$, (b) $p = 0.25$.

In order to evaluate the impact of adding a loss detection service, we consider two protocols noted A and D . Both of them benefit from NACK suppression and subcast services but D , in addition, benefits from a loss detection service at the active routers (see section 2.3.6).

In order to consider the detection-capable active router's position, we defined T_{ar} (respectively T_{sr}) as the number of links between the router (respectively the source) and any receiver. We then introduced $\alpha \in]0, 1[$ which is the ratio of T_{ar} to T_{sr} . We considered two particular values of α : $\alpha_c = (T_{sr} - 1)/T_{sr}$ where the router is one link far from the source ($T_{sr} - 1$ links between the router and a receiver) and $\alpha_f = 1/T_{sr}$ where the router is one link far from a receiver. Figure 8 shows the performances of D as a function of the router position. It plots the ratio of the achieved delay by D when we change the router position (parameter α) to the case where $\alpha = \alpha_f$. We can see that the position of the router with respect to the source is deciding. The gain is increased when the router is the closest to the source. The gain in delay observed is more important for high loss rates. In fact, for a loss rate of 50% the delay is improved by 25% up to 35% if we move the router near the source instead of putting it near the receivers. As a conclusion one must be careful about where to place a loss detection capable-router. This latter must neither be too far from nor too close to the source. When the router is closer to the receivers, the load introduced by the loss detection service is unjustified because of the long distance to be crossed by the generated NACKs by the router. When the router is put sufficiently far from the source, we maximize the number of losses that can be detected. When the position of the active router is well chosen, we showed that D performs better than A especially for high loss rates.

A set of simulations are performed to show how a loss detection service could decrease the delay of recovery of the DyRAM framework. Let $DyRAM^+$ be the enhanced version of $DyRAM$ resulting from adding the loss detection capability. To do so, the four protocols A , D , $DyRAM$ and $DyRAM^+$ are simulated on a network model depicted in figure 9 similar to the adopted one (Figure 3) except that a router close to the source (the source router A_S) is considered. All the routers are assumed to perform at least two active services, the NACK suppression service and the subcast functionality. Based on our analytical results, the source router is the best candidate to perform the loss detection service in addition to the two previous services. The other active routers should only perform the replier election service as seen in the previous section.

Figure 10 plots the recovery delay (normalized to the RTT) for the four protocols as a function of the number of the receivers for different loss rates. First of all, it is noticeable that protocol $DyRAM$ and $DyRAM^+$ (both with local recovery from the receivers) always perform better. For instance, we can see in Figure 10 that $DyRAM$ is 3 times faster than A for a loss rate of 25%. Now, when the loss detection service is applied to A (giving protocol D) the recovery delay can be reduced. In fact as we can see, D performs better than A thanks to the loss detection service. When applying the loss detection service to




Figure 8. The delay gain in D as a function of the router position.




Figure 9. Network model.

DyRAM, the delay of recovery decreased mainly for high loss rates and a large number of receivers. For instance, the loss detection service allows *DyRAM* to go 4 times faster for 96 receivers and a loss rate of 25%.




Figure 10. The recovery delay with $p = 0.25$.




Figure 11. An active-based reliable multicast architecture

4. An Active-based Reliable Multicast Architecture

We consider the IP multicast model introduced by Deering [3]. The source sends messages to a multicast address subscribed to by all the receivers. The sender does not need to know the receivers identities. A receiver can join or leave a multicast session as it wishes by simply sending a *join* or a *leave* message via IGMP (Internet Group Management Protocol). The routers has to make their best effort to deliver the traffic from the sender to each receiver. The packets will be duplicated in the network as needed, thus reducing bandwidth consumption compared to the unicast approach. Our reliable protocol DyRAM works on top of UDP and will guarantee that all the sent packets are received by all the receivers of the multicast group.

Our architecture (see Figure 11) is based on a virtual topology of active routers which are able to perform customized computations (services) on the messages flowing through them. We only consider active routers at the edge of the core network. This is due to the fact that the core network is reliable and a very high-speed network. Adding complex processing functions inside the core network will certainly slow down the packet forwarding functions.

All the active routers are assumed to perform at least two active services, the NACK suppression service and the subcast functionality. The active router which is located at the sender's side just before the core network is called the *source router*. As have been shown, the loss detection service is only beneficial if the loss detection capable-router is close enough to the source. Consequently the source router is the best candidate to perform the loss detection service in addition to the two previous services. The other active routers should perform the replier election service. This is essential to avoid electing as a replier, a receiver which is located very far from the affected subtree.

For the realization of our active approach, we propose the *programmable switch approach* which maintains the existing packet format. Programs are injected separately from the processing of messages. Initially, and using a reliable channel, the source injects the required service by sending them to the multicast address. Afterwards, the source begins the multicast of the data packets. When an active router receives a packet, it first looks for the appropriate service deduced from the packet header. Consequently a recognized packet as a DyRAM one will be processed by the DyRAM active services that process in a customized way the different packet types.

To dynamically handle the multicast group topology changes, active routers have to be able to add or remove the concerned services. An active router that leaves the multicast session has simply to remove the associated services. However when an active router joins a multicast tree, it has to easily download the required services. This can be achieved by requesting the services from the closest active router which has already installed them. If there is no such active router then the services need to be sent from the source. The required services are always transmitted via a reliable channel. These operations are usually handled by the used active execution environment.

5. Implementation and Experimental Measurements

The performances of DyRAM have been extensively studied in the previous section by simulation. Results show that local recovery associated to early loss detection succeed in providing low latencies. However, given the reliance of DyRAM on active services, it is crucial to determine how costly these services are. In this section we report our experience in prototyping DyRAM on top of the TAMANOIR execution environment [6]. At this point, we must mention that the purpose of this work is to quantify the raw processing time required for each elementary service and that neither timer optimizations nor end-to-end recovery latencies are investigated.

5.1. AN IP ADDRESS-BASED REPLIER ELECTION

DyRAM is initially intended to rely on a network multicast facility such as IP multicast. However, as an alternative to IP multicast, it is possible to have active services realizing level 3 multicast functionalities. Once a group has been formed relying on IGMP for instance, DyRAM runs its own simple session protocol to gather additional topological information at the DyRAM level to enhance the group anonymity imposed by IP multicast. A router mainly gathers information on the IP addresses of receivers or active routers directly connected to it.

For our first DyRAM prototype, we adopted an implementation based on the use of IP addresses instead of link information. We need to gather some information at the beginning of the multicast session. Consequently an initialization phase is required. At the beginning of the multicast session, the source sends to the the hole group an initialization message (INIT). Upon the reception of this message, a receiver will respond with an INIT_RESP message that contains its IP address. The first router that receives an INIT message will mark it as already received by at least one active router of the multicast tree. This first crossed router will behave as a source router with the loss detection facility as proposed in section 4. Moreover such a router will forward immediately a NACK to the source instead of wasting time in a replier election procedure. Each router of the first level (just upstream a receiver) retrieves the IP address of the downstream receivers from their INIT_RESP messages. Afterwards, a first level router will forward one INIT_RESP message while the others will be dropped. An i th level router (with $i \geq 2$) retrieves one IP address per downstream link from the INIT_RESP messages it receives.

In such a manner, the source would receive just one INIT_RESP message which will be dropped. A timer whose the expiration triggers the emission of the HB_RESP_S message (described in section 2.5), could be initialized. Motivations behind waiting for a given amount of time is to give time to the other nodes to have a first estimation of their RTTs to their parents. INIT and INIT_RESP messages may be lost. This is why the source, when it sends an INIT message, sets a timer. The expiration of this timer without the reception of the INIT_RESP message, triggers the emission of a new INIT message. CR messages can also used to update periodically the IP addresses kept in a router for the purpose of the replier election. This is done to quickly react to eventual topology changes.




Figure 12. Used topologies (a) DP, NACK and DR service, (b) replier election

5.2. PACKET FORMAT AND DATA STRUCTURES

DyRAM uses the ANEP format for the packets which is the one handled by TAMANOIR. In the ANEP header appears the multicast and the source addresses in addition to the service identifier. The other fields are inserted in the payload. DyRAM is implemented on top of the UDP protocol since the reliability is insured by DyRAM itself. Incoming DyRAM packets are given to the TAMANOIR execution environment which selects the appropriate active service based on the *SVC* (service identifier) field.

In this first DyRAM version, the TL structure is implemented as a double-linked list of bits (based on an array implementation) that allows fast insertion/suppression of bits at the beginning/end of the list. The NS structure is handled with a hash table with the sequence number of the lost data packet as the hash key. For this first prototype, we chose to use IP addresses of the receivers instead of the links information. Each entry refers to a double-linked list (*subList*) of IP addresses. In the current implementation, timers are handled by threads and a hash table of timer threads is used to maintain the list of timers. We will see later on that it may be not the best design choice.

5.3. TEST-BED AND SCENARI II

Experimentations were performed with a first prototype using a test-bed that consists in a set of receivers and 2 PC-based routers (Pentium II 400MHz, 512KByte cache, 128MByte RAM) running a Linux 2.4 kernel and Java version 1.3.1. The purpose is the measurement of the overhead incurred by the active services, a set of experiments were performed in two different configurations. In the first configuration (see Figure 12a), the source *ike* multicasts packets to *stan* and *resamd* through 2 active routers (*resama* and *resamo*). In order to measure the data service overheads in this topology, data packet losses are produced between the source and the first active router *resamo*. Figure 12a shows the steps of the recovery process where the source is the replier: DP and DR refer respectively to Data Packet and Data Repair. In the second configuration (Figure 12b), the source *ike* multicasts packets to a set of receivers through the active router *resamo*. The cost of the replier election is mainly the processing time for comparing the *subList* list against the receiver list in order to determine a replier for the lost data packet.

5.4. COST OF ACTIVE SERVICES

The cost of the data packet service represents the processing time required to forward the packet to the destinations using the underlying IP multicast functionalities. Therefore, this cost mainly includes the update of the track list when there are no sequence gaps. In case of a gap sequence indicating a packet loss, the data service includes additionally the spent time in setting a timer for ignoring subsequent similar NACKs. The cost of the NACK service contains the processing time to update or create the NS

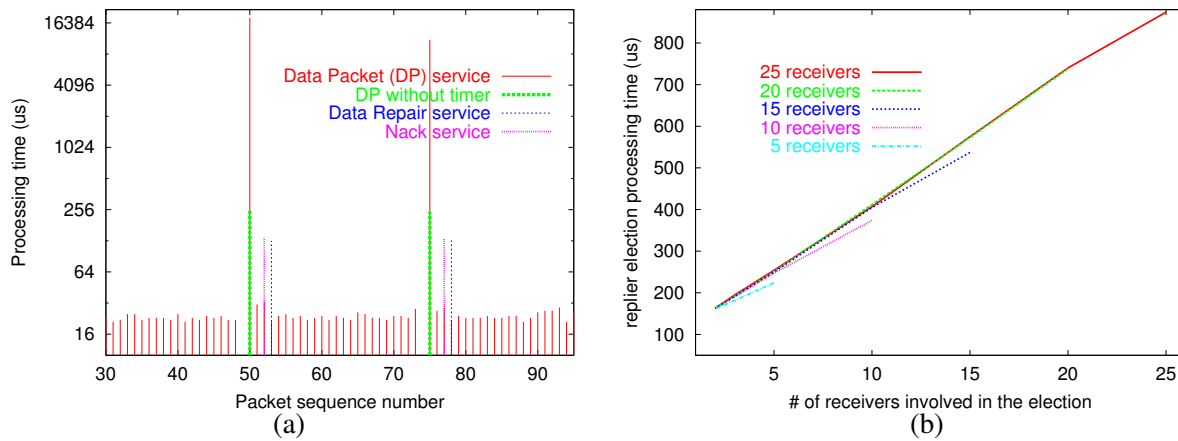


Figure 13. Cost of (a) packet services, (b) replier election.

structure. The cost of the data repair service represents the processing time of scanning the *subList* and performing the subcast. Preliminary results on our test-bed are illustrated in Figure 13. Figure 13a shows the processing time in μs of a data packet, a repair packet and a NACK packet at the *resamo* active router. The x-axis shows the packet sequence number while the y-axis shows the processing time. The data packet size is initially set to 4KByte.

We can see that, in the absence of packet losses, the processing time of a data packet is about $20\mu\text{s}$. A gap in the x-axis represents a packet loss. For instance, in our test scenario, packet 49 is first lost thus making the processing of packet 50 longer because of the loss detection and associated data structure updates (track list and NS structures). Figure 13a shows that each packet loss incurs an additional cost of about 12ms to 17ms for the next packet that corresponds to the processing time for the data packet and the time for setting up a timer thread for NACK discarding. It is worth mentioning that the first overhead is only about $250\mu\text{s}$! Several optimizations are possible for the timer management (reuse of old threads or the use of only one thread with a hash table for all timers for instance.) and we expect much lower overheads in next implementations. The NACK packet and the repair packet that follow the lost packet are processed in approximately $135\mu\text{s}$ and $123\mu\text{s}$ respectively (for these packets, the x-axis indicates the current data packet sequence number at the moment they are processed in order to respect the chronological order).

In this first prototype, the election is performed on-the-fly. A default replier is selected when the first NACK arrives and the choice of the replier is updated at each NACK reception. The advantage of this approach is to perform most of the election processing within the timer duration for gathering NACK information resulting in no cost from a protocol perspective. Figure 13b shows the processing time to dynamically determine the replier. The number of receivers under the active router ranges from 5 to 25. The x-axis indicates how many receivers are involved in the replier election process. For instance, if we look at the 5-receiver curve, a value of 5 at the x-axis means that a replier is found after having scanned 4 receivers. The curves show that the number of total receivers has little impact of the election process: finding a replier after 4 trials amongst 25 receivers adds $30\mu\text{s}$ to the case where the replier is found amongst 5 receivers.

6. Conclusion

In this paper we have reported our experience on the design, evaluation and the implementation of DyRAM, an active reliable multicast protocol. DyRAM performs the local recovery from the receivers with the assistance of the routers without any in-network caching facilities. A link is elected as a replier on a per-packet basis to perform the retransmission of a data packet. An estimation of the link capacities

is performed in order to have a fairer replier election with a nice load balancing feature among the links downstream a router.

Knowing that losses are mainly due to congestion, a congestion control algorithm is already proposed to strengthen our reliability mechanism. In this paper we focused on the reliability mechanism with a brief overview of the congestion control algorithm. However an extensive evaluation study using both simulations and analysis has been performed on DyRAM. In addition a first prototype of DyRAM has been presented with preliminary measurements of the raw processing time required by the different services.

Regarding the problem of scalability, our approach does not use caching facilities within routers but only a few amount of RAM memory for maintaining some data structures per multicast session. As the amount of memory needed for this purpose is small (few KBytes per session) we believe our approach to be more scalable than an approach like ARM. Based on small data structures we showed via the implementation of DyRAM that its associated active services introduce only low overheads in routers and that implementations without much performance degradation on the routing and forwarding functions are possible. In fact in this report we addressed this performance issue by measuring the per-router processing cost of elementary services proposed in the DyRAM protocol. The results are very encouraging as most processing costs range from tens to hundreds of microseconds on a Pentium II 400MHz PC-based router.

Given these results, we believe it is very possible to build active routers from regular PCs (using the most up-to-date processor and clock rate) and still get performances at high bit rates. This possibility would certainly help to disseminate the use of active networking technologies for a large range of applications as the deployment would be easier and faster than a dedicated router solution.

References

1. R. Bagrodia et al. Parsec: A parallel simulation environment for complex systems. *Computer Magazine*, 1998.
2. M. Calderón, M. Sedano, A. Azcorra, and C. Alonso. Active networks support for multicast applications. *IEEE Networks*, May/June 1998.
3. Steve. Deering. Distance vector multicast routing protocol. RFC 1112, IETF, August 1989.
4. S. Zabele et al. Improving distributed simulation performance using active networks. In *World Multi Conference*, 2000.
5. Sally Floyd, Van Jacobson, Ching-Gung Liu, Steven McCanne, and Lixia Zhang. A reliable multicast framework for light-weight sessions and application level framing. *IEEE/ACM Transactions on Networking*, 5(6), 1997.
6. J.P. Gelas and L. Lefèvre. Tamanoir : A high performance active network framework. Workshop on Active Middleware Services 2000, 9th IEEE International HPDC, Pittsburgh.
7. Hugh W. Holbrook, Sandeep K. Singhal, and David R. Cheriton. Log-based receiver-reliable multicast for distributed interactive simulation. In *SIGCOMM*, Oct. 1995.
8. S. Kasera and S. Bhattacharya. Scalable fair reliable multicast using active services. *IEEE Network Magazine's Special Issue on Multicast*, 2000.
9. L. Lehman, S. Garland, and D. Tenenhouse. Active reliable multicast. In *Proc. of the IEEE INFOCOM, San Francisco, CA*, March 1998.
10. Moufida Maimour and Cong-Duc Pham. An analysis of a router-based loss detection service for active reliable multicast protocols. In *Proceedings of the 11th IEEE International Conference on Networks*, Singapur, August 2002.
11. Moufida Maimour and Cong-Duc Pham. Dyram : A reliable multicast protocol. Technical Report RR-4635, INRIA, 2002. Also available as a LIP/ENS Research Report under 2003-05.
12. Moufida Maimour and Cong-Duc Pham. Amca: an active-based multicast congestion avoidance algorithm. In *the 8th IEEE Symposium on Computers and Communications (ISCC 2003)*, pages 747–754, Kemer-Antalya, Turkey, July 2003.
13. Christos Papadopoulos, Guru M. Parulkar, and George Varghese. An error control scheme for large-scale multicast applications. In *Proc. of the IEEE INFOCOM*, March 1998.
14. S. Paul and K. Sabnani. Reliable multicast transport protocol (RMTP). *IEEE JSAC, Spec. Issue on Network Support for Multipoint Communications*, 15(3), April 1997.
15. T. Speakman, J. Crowcroft, J. Gemmell, D. Farinacci, S. Lin, D. Leshchiner, M. Luby, T. Montgomery, L. Rizzo, A. Tweedly and N. Bhaskar, R. Edmonstone, R. Sumanasekera, and L. Vicisano. Pgm reliable transport protocol specification. RFC 3208, IETF, December 2001. Appeared also in *IEEE Networks Magazine, Special issue on "Multicasting: An Enabling Technology*, 2003.

16. R. State, O. Festor, and E. Nataf. A programmable network based approach for managing dynamic virtual private networks. In *Proceedings of PDPTA'02, Las Vegas, June 26-29, 2000*.
17. D. L. Tenenhouse, J. M. Smith, W. D. Sincoskie, D. J. Wetherall, and G. J. Winden. A survey of active network research. *IEEE Communication Magazine*, pages 80–86, January 1997.
18. Rajendra Yavatkar, James Griffioen, and Madhu Sudan. A reliable dissemination protocol for interactive collaborative applications. In *ACM Multimedia*, 1995.

Address for Offprints:

Ecole Normale Supérieure de Lyon
46, Allée d'Italie
69364 LYON Cedex 07
Tel : +33 (0) 4 7272 8759
fax : +33 (0) 4 7272 8080

