



HAL
open science

Performance Evaluation of DyRAM using NS.

Moufida Maimour

► **To cite this version:**

Moufida Maimour. Performance Evaluation of DyRAM using NS.. [Research Report] LIP RR-2004-12, Laboratoire de l'informatique du parallélisme. 2004, 2+16p. hal-02102000

HAL Id: hal-02102000

<https://hal-lara.archives-ouvertes.fr/hal-02102000v1>

Submitted on 17 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Laboratoire de l'Informatique du Parallélisme

École Normale Supérieure de Lyon
Unité Mixte de Recherche CNRS-INRIA-ENS LYON-UCBL n° 5668

***Performance Evaluation of DyRAM using
NS***

Moufida Maimour

March 2004

Research Report N° 2004-12

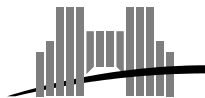
École Normale Supérieure de Lyon

46 Allée d'Italie, 69364 Lyon Cedex 07, France

Téléphone : +33(0)4.72.72.80.37

Télécopieur : +33(0)4.72.72.80.80

Adresse électronique : lip@ens-lyon.fr



Performance Evaluation of DyRAM using NS

Moufida Maimour

March 2004

Abstract

Since it is required by many emerging Internet applications, reliable multicast has received significant attention in the networking research community. Scalability to a large number of receivers is one of the key challenges to be met by reliable multicast protocols. We previously proposed a reliable multicast protocol called DyRAM that benefits from the assistance of routers in order to implement its services. In this report, we consider a simulation-based performance evaluation of DyRAM. We mainly compare it to SRM and a native reliable multicast protocol. We show that DyRAM allows for less load in terms of consumed bandwidth on the different multicast nodes and that the different services allows for reducing the repair latency and implosion. We also consider the active routers density and show that increasing the density of active routers allows for better performances. Comparing DyRAM to SRM allows for showing the load balancing feature and scalability of our approach.

Keywords: Reliable multicast, Simulation, Performance evaluation, ns

Résumé

Comme le multicast fiable est requis par un nombre croissant d' applications, il a reçu un intérêt significatif de la part de la communauté de recherche en réseaux. Le support d'un grand nombre de récepteurs est l'un des défis majeurs des protocoles de multicast fiable. Dans le passé, nous avons proposé un protocole de multicast fiable appelé DyRAM qui bénéficie de l'assistance des routeurs afin d'implémenter ses services. Dans ce rapport, nous présentons une évaluation des performances de DyRAM par simulation. Principalement, nous le comparons à SRM et un protocole de multicast fiable natif. Nous montrons ainsi que DyRAM permet de diminuer en termes de bande passante consommée, la charge des différents nœuds et que ses services permettent la réduction du temps de recouvrement ainsi que le problème d'implosion. Nous considérons également la densité des routeurs actifs et nous montrons que son augmentation améliore les performances. La comparaison de DyRAM à SRM montre la distribution équilibrée des charges ainsi que le passage à l'échelle de notre approche.

Mots-clés: Multicast fiable, simulation, évaluation de performances, ns

1 Introduction

Since it is required by many emerging Internet applications, reliable multicast has received significant attention in the networking research community. Scalability to a large number of receivers is one of the key challenges to be met by reliable multicast protocols. Tree-based approaches have been proven to be the most efficient. Organizing the receivers in a hierarchy allows for the implementation of local recovery, thus unloading the source from performing retransmissions in addition to reducing the recovery latency. Moreover, receivers' feedback can be suppressed through the hierarchy, thus avoiding feedback implosion. Finally, retransmissions can be more localized thanks to the hierarchy structure, thus eliminating duplicate and unwanted repairs.

In a tree-based approach, mechanisms to build and maintain the hierarchy structure are required. To be more efficient, the built hierarchy has to be congruent with the underlying multicast distribution tree in addition to its ability to adapt to the dynamic nature of a multicast group. DyRAM (for *Dynamic Replier-based Active reliable Multicast*, our reliable multicast protocol [4]) is one protocol that satisfies both congruence and adaptability requirements. In DyRAM, intermediate nodes (routers) in the multicast tree are involved in the data recovery process. On one hand, a router contributes in the election of "repliers" to retransmit lost packets, thus building an implicit recovery tree that is congruent with the physical multicast tree. On the other hand, a repplier is elected on a per packet-basis thus allowing for an adaptability to dynamic changes in the multicast group topology. Since our proposal is router-assisted, additional design goals are required. Mainly, we have been careful about reducing the amount of buffering requirements at the network elements and minimizing the overhead at the routers by implementing light-weight functionalities.

In this report, we investigate an ns-based evaluation of DyRAM. Ns has gained wide acceptance among the network community. For the purpose of the performance evaluation of DyRAM (only the reliability mechanism without any congestion control), the ns model has allowed for a comparison study with another reliable multicast protocol SRM [2]. SRM implements the local recovery from the receivers using a timer-based approach. It is worth mentioning that other reliable multicast protocols are also implemented using ns, however we were not able to perform comparisons with them. This is due to the fact that they either are not or no longer available on the Internet like AER [3] or do not integrate all the functionalities like PGM [8] whose the implemented model does not yet integrate the recovery from the receivers.

The remainder of this report is organized as follows. First, an overview of DyRAM and SRM is given in section 2. Afterwards, the simulation model is presented in section 3. The simulation results are presented in 4. Many aspects are dealt with such as the impact of the routers density, the load balancing at the receivers and the number of the receivers subscribed to the multicast session. Finally, section 5 concludes.

2 Background

2.1 DyRAM Overview

DyRAM is a reliable multicast protocol with a recovery strategy based on a tree structure constructed on a per-packet basis with the assistance of the routers. DyRAM 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 repplier election, positive acknowledgements (ACKs) are piggybacked on negative ones (NACKs). In the absence of NACKs, ACKs are also periodically piggybacked on special messages called "Congestion Reports"(CRs). These CRs contain the necessary information for the flow and congestion 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 :

- the NACK suppression of duplicate NACKs in order to limit the NACK implosion problem.

- the CRs aggregation thus avoiding the CRs implosion at intermediate nodes and at the source.
- the 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 the affected subtree.
- the 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 in addition to load-balancing features.
- the early loss detection of packet losses and the emission of the corresponding NACKs. This is very helpful in providing a low recovery delay.

2.2 SRM Overview

Scalable Reliable Multicast (SRM) [2] designed to assure reliability to a white board (wb) application on the MBONE, is one of the first proposed protocols that exploit local recovery from the receivers. Upon the reception of a loss, a request is sent on the multicast address. Any subscriber to the group could process this request if it has already received the corresponding data packet. In order to avoid that all the receivers be flooded by multiple NACKs and/or duplicate repairs, a probabilistic approach is adopted to perform NACK/repair suppression. This approach is similar with some enhancements, to the one originally proposed by [9] in the context of the *eXpress Transport Protocol* (XTP). Upon the detection of a loss, a receiver waits for a random period determined by its distance from the original source of data before sending its NACK. In such a way, the closest host to the point of failure is likely to timeout first and multicast the request even if a number of hosts may miss the same packet. These hosts will hear that request and suppress their own and behave as they have sent the corresponding NACK. Similarly and upon the reception of a request, a receiver (if it has already received that corresponding packet) will set a repair timer to a random value depending on its distance from the request originator. A host close to the affected area is likely to timeout first and send the corresponding repair. This timer will be canceled if the repair is received; otherwise the requested packet is retransmitted on timer expiration.

A NACK-based approach do not allow the sender to free its buffers kept for potential retransmissions. This is why a combined ACK/NACK approach could be a better solution when combined with a structure-based approach like RMTP [7], TMTP [10] or TRAM [1] (each of them will be described below). Organizing the receivers in a logical structure is mainly adopted to perform feedback aggregation/suppression, thus avoiding the feedback implosion problem. Moreover it provides a better control regarding the request/repair exposure. In SRM, a TTL-based approach has been proposed to limit the scope of NACKs and retransmissions, however this still an immature solution. How the TTL field is set is not an easy task since for instance, a request could not reach a potential replier. In structure-based approaches, requests are sent to a dedicated receiver instead of multicasting them to the whole group. Thus avoiding the *crying-baby* syndrome that can be observed in a SRM session where a weak receiver is overwhelming the others with very frequent requests.

3 Simulation Model

3.1 Network and Loss Model

We consider a packet switching network composed of a fast core network and several slower edge access networks. One source multicasts data packets to R receivers distributed among N local groups. A local group is composed of the set of receivers $\{R_{i1}, R_{i2}, \dots, R_{iB}\}$ located downstream a router A_i , $i = 1..N$ (figure 1). Active routers are only considered at the edge of the core network. The core network is reliable and a very high-speed network and adding complex processing functions inside the core network will certainly slow down the packet forwarding functions. The router

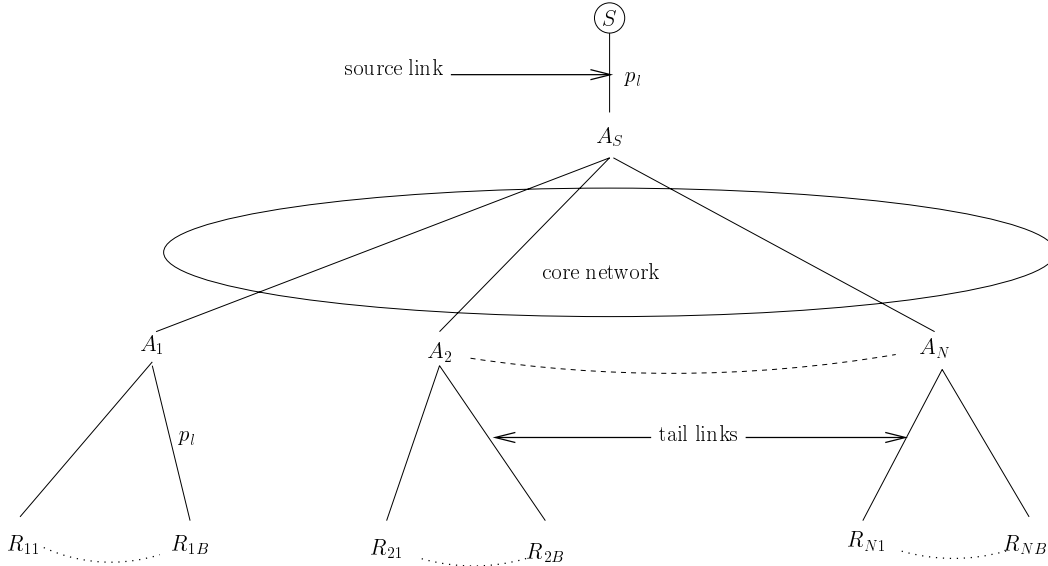


Figure 1: Network model.

located at the source side (A_S) is called the *source router* and the others located at the receivers side are designated by *tail routers*.

The number of the receivers per each group is chosen randomly following a uniform distribution of parameters B_{min} and B_{max} . Furthermore, the bandwidth and the delays associated to each link are also chosen randomly. In such a way, the receivers of a local group are not of equal distance from their active router and have different capacities. In this network model, a router can be configured as active or passive. This is interesting when considering the impact of active routers density on the performances of DyRAM.

In order to generate losses at the different lossy links, a loss module is assigned to every tail and source link with a loss rate of p_l . We consider, three different patterns of losses. The first pattern consists in correlated losses introduced at the source link. The second pattern produces losses only at the tail link. The third pattern combines the two previous patterns where losses occur at both the source link and the tail links. The transmission rate of the source is fixed appropriately so losses due to congestion are avoided. The congestion avoidance algorithm associated to DyRAM is not considered here.

3.2 Protocols and Scenarii

In addition to DyRAM and SRM, we also consider a native reliable multicast protocol noted S which does not benefit from active processing at the routers. In S , there is no local recovery, all the retransmissions are performed by the source and additionally there is no feedback suppression at the routers.

For the case of DyRAM, all the active routers are assumed to perform at least two active services, the NACK suppression service and the subcast functionality. The replier election is only performed by the tail routers while the loss detection is only performed by the source router according to our architecture proposed in [6]. In all the simulations, the replier election is performed with respect to a ring order without any optimization.

All the simulations unless stated otherwise, are run for 200 seconds with a multicast file of an unlimited size. Losses are introduced at 10s and are stopped at 110s. B_{min} and B_{max} are set respectively to 1 and 20.

3.3 Metrics

The metrics used for the evaluation are the following:

Repair Latency The repair latency for a lost packet is the time elapsed between the emission of a request and the reception of the first repair normalized to the RTT between the affected receiver and the source. In our results, a mean of the repair latency is computed among all the losses experienced by all the receivers.

Load at the Different Nodes For each type of node (source, receiver or router), the amount of input and output bytes per unit of time is recorded during the simulation. This gives an idea on the load experienced by each node in the network model. This is useful for instance, to have an idea on the source load and the contribution of the receivers in the recovery process.

Repair Implosion The repair implosion is quantified by computing a mean among the duplicate repairs per packet obtained by dividing the number of repairs received by all the receivers on the number of NACKs sent for this data packet.

4 Simulation Results

4.1 Benefit of Active Routers

The purpose of this section is to show quantitatively the benefit of the existence of active routers in DyRAM. To do so, simulations are conducted on DyRAM and S during 200 seconds for different loss scenarii.

4.1.1 Feedback Suppression

Figure 2 plots for a loss rate of 10% per link, the input load (corresponding to the feedback messages) at the source where the losses occur only at the source link. In this case, there is no local recovery since a loss affects all the subscribed receivers. We can see that the source in S receives more feedback than in DyRAM. In the absence of losses (before 10s and after 110s). This is due to the fact that the CRs sent periodically by the receivers in S are not aggregated by the routers before reaching the source. When losses occur (between 10s and 110s), we note that the amount of feedback received by the source increases in both S and DyRAM due to the transmission of NACKs by the affected receivers. We can deduce from figure 2, that the source receives less NACKs in DyRAM compared to S . This reflects the NACK suppression performed by DyRAM at the active routers before they reach the source. In summary, the source is unloaded in DyRAM from processing multiple feedback sent by the receivers. For instance, the amount of feedback is reduced by a factor of approximately 4 and 3 in respectively the presence and the absence of losses.

4.1.2 Local Recovery

Figure 3 shows the input load at the source when losses occur at the tail links. This figure shows the benefit of the loss recovery from the receivers instead of the source. Since all the losses are experienced at the tail links, most of the losses in DyRAM can be recovered locally. Indeed, the source in DyRAM as we can see in figure 3, is unloaded during all the simulation. However, in S the source is overwhelmed by feedback messages that mainly consist in requests from the receivers affected by losses. These feedback are not suppressed at the routers as in DyRAM. We can see that the source examines a load multiplied by a factor of 90 compared to the case of DyRAM.

Figure 4 shows the amount of feedback the source receives when losses occur at both the source and tail links. Once again, the source is unloaded in DyRAM thanks to the active services such as CR aggregation, NACK suppression and replier election. Furthermore, we can notice that all




Figure 2: Feedback at the source in S and DyRAM, $p_l = 0.10$ with losses at the source link.




Figure 3: Feedback at the source in S and DyRAM, $p_l = 0.10$ with losses at the tail links.




Figure 4: Feedback at the source in S and DyRAM, $p_l = 0.10$ with losses at all the links.




Figure 5: Output load at a receiver in S and DyRAM, $p_l = 0.25$ with losses at (a) the source link, (b) all the links.

the losses that occurred are recovered earlier in DyRAM compared to S . Losses are recovered at 115s in DyRAM instead of 180s in protocol S .

4.2 Load at the Receivers

Figures 5 and 6 plot respectively, the evolution of the output and input load at a receiver in DyRAM and S for a link loss rate of 25%. When losses occur at the source link, we can see in figure 5a that the two protocols are equivalent in terms of the amount of feedback sent by the receivers. In such a case there is no local recovery from the receivers which in both S and DyRAM, have similar behavior when sending CRs and NACKs to the source. When losses occur at all the links, we can see in figure 5b that a receiver in DyRAM sends out more bytes in the period of losses. This additional traffic from a receiver in DyRAM consists in the repairs sent by this receiver. However, this additional traffic is no longer present at the receivers side. Thanks to the local recovery, losses are rapidly recovered. In S , we can observe that all the losses are not recovered even by the end of the simulation.

With respect to figure 6a, we can see that once again, S and DyRAM have similar behavior in terms of the input load at the receivers when the losses occur at the source link. However, when losses occur also at the tail links as shown in figure 6b, a receiver in S has no longer the same behavior as in DyRAM. We can see that after the loss period, there is more load at the receivers in DyRAM since they continue receiving repairs and NACKs in addition to original data packets




Figure 6: Input load at a receiver in S and DyRAM, $p_l = 0.25$ with losses at (a) the source link, (b) all the links.

without losses. This will no longer be the case and losses are recovered more quickly in DyRAM than in S . In this latter, we can observe that a receiver continues receiving repairs until the end of the simulation.

4.3 Active Routers Density

In this section, we consider the density of the active routers (those located at the receivers' side) on the performance of DyRAM. In order to be able to have the largest multicast group, the number of data packets sent by the source is limited to 200 packets instead of the unlimited number for the previous simulations. Losses occur during all the simulation time and the simulation ends when all the receivers have received all the 200 packets sent by the source. The simulations are run for different values of $N = 1, 2, 3, 4, 8, 20, 40$. The number of the receivers per router is chosen randomly from a uniform distribution of parameters 1 and 20. A number of simulations is performed and a mean is computed for experiences with the same number of receivers.

Figure 7 plots for different loss rates and scenarii, the repair latency of lost data packets for different densities (0, 50 and 100%) of active routers. First, for low loss rates say 1% as depicted in figures 7a1-a2, we note that when the number of the receivers is small, there is no noticeable difference between an active (with 50 or even 100% of active routers) and a traditional approach in terms of repair latency. When the number of the receivers increases, we note the increasing benefit of the existence of active routers in the multicast tree. For instance, for more than 400 receivers, we can examine a repair latency reduced with a factor of 1.8 when 50% of the routers are active compared to the passive approach (protocol S). When all the routers are active, the repair latency is improved with a factor 2.5 and 4.25 compared to the case of 50% of active routers and the passive approach respectively. When the loss rate increases (figures 7b-c) and especially in the presence of a large number of receivers, we can easily note the benefit of increasing the number of active routers in a reliable multicast session. We observe a repair latency reduced by a factor 2 to 5 when a passive approach (S) is replaced by a fully active approach (DyRAM with 100% of active routers).

4.4 Comparing DyRAM to SRM

DyRAM implements local recovery from the receivers instead of overwhelming the source by multiple retransmissions. SRM[2] is one of the first protocols that moved the responsibility of loss recovery to the receivers. This section is an attempt to compare our approach DyRAM to SRM mainly, in terms of the load at both the routers and the receivers. We mainly consider the load balancing at the receivers and the impact of the number of the receivers on the performances. In




Figure 7: Influence of the active routers density in DyRAM on the repair latency, (a) $p_l = 0.01$, (b) $p_l = 0.05$, (c) $p_l = 0.1$. Losses at the tail links (left side) and losses at all the links (right side).




Figure 8: Consumed bandwidth by a tail router in SRM and DyRAM, $p_l = 0.1$.

what follows, we only consider losses at the tail links since the purpose is to evaluate the local recovery from the receivers instead of the source.

4.5 Load at the tail routers

Figure 8 plots the bandwidth consumed by a tail router in DyRAM and SRM when the loss rate per link is 0.1. We can see that in the absence of losses (before 10s), the amount of consumed bandwidth by a tail router in SRM and DyRAM is almost constant. We observe that SRM has more load at the tail routers than DyRAM. This is due to the amount of signaling messages generated within SRM. When losses occur, a tail router becomes very loaded in SRM compared to DyRAM. In SRM, NACKs and repairs are sent on the multicast address to all the participants and hence received by all the routers of the multicast tree. In DyRAM, a NACK is sent toward the source and is intercepted by the routers located in its path. These routers process NACKs they receive in order to suppress duplicate ones. Moreover, in DyRAM, a repair from the source or any receiver is sent only on the affected links of the multicast tree thanks to the subcast mechanism. In SRM, the additional load at the routers is due to the reception of the different NACKs sent by all the affected receivers and the multiple receptions of repairs for the same loss.

4.6 Load at the receivers

Figure 9 shows the input load at a receiver for a loss rate of 25% per link. The input load of a receiver is the amount of data per second it receives during the multicast session. Mainly, a receiver receives repairs in addition to original transmissions. Due to local recovery from the receivers, a receiver can receive requests (NACKs) from other receivers in the multicast session. We notice in figure 9, that in the absence of losses, a receiver in DyRAM would have similar input load as in SRM. However when losses occur, an SRM receiver is more loaded than a DyRAM one, except for a short period that follows the losses interval as can be observed in figure 9. This shows the earlier reception of repairs by a receiver in DyRAM while it continues the reception of original transmissions without losses. An SRM receiver receives more requests compared to DyRAM because of duplicate NACKs received from the other members of the group. The load at a receiver in SRM is not reduced even at the end of the simulation. This means that losses are not completely recovered.

Figure 10 shows the output load in DyRAM and SRM at a receiver for a loss rate of 25% per link. The output load is the amount of bytes sent by a receiver in the two considered protocols. A receiver sends feedback messages (mainly NACKs) and to implement local recovery, can send repairs for NACKs it receives. Once again, we notice that in the absence of losses, a receiver in DyRAM would have similar input load as in SRM. During the lossy period, the two protocols are




Figure 9: The input load at a receiver in SRM and DyRAM, $p_l = 0.25$.




Figure 10: The output load at a receiver in SRM and DyRAM , $p_l = 0.25$.

equivalent in terms of the amount of bytes sent by a receiver. Each receiver sends requests for missing data packets in addition to transmitting repairs when requested by other receivers. We can also observe that losses are recovered much earlier in DyRAM. In SRM, a receiver continues to be overloaded by sending requests for packets it has not received yet in addition to retransmitting repairs for duplicate NACKs received from the other members of the group.

4.7 Load Balancing at the Receivers

In DyRAM, a per-packet replier election is performed by the routers. This is a nice feature, since a load balancing among the receivers of the same local group, is possible. Here, we aim to show how the retransmission load is distributed among the receivers in DyRAM. Figure 11 plots for DyRAM (left side) and SRM (right side), the input, output and overall load at three different receivers for a loss rate $p_l = 0.01$. First, we note that there is more load on the receivers in the lossy interval (10s-110s). In terms of load balancing at the receivers, we observe that in both of the two protocols, the different receivers have the same amount of data coming in or going out. This shows that for a low loss rate, the two compared protocols (SRM and DyRAM) are equivalent in terms of load balancing among the receivers.

To see the impact of higher loss rates on the load balancing in the two protocols, figure 12 plots for a loss rate of 25% per link instead of 1%, the input, output and overall load at three different receivers. We observe that in DyRAM (left side), the different loads are the same at the three receivers. However, we observe in SRM (right side) that the three considered receivers

present different loads. As we can see in figure 12a2, there is one node that receives more NACKs (about 12.5% of additional NACKs) and thus retransmits more repairs (about 100% of additional repairs) than the other nodes as depicted in figure 12b2.

4.8 Scalability Evaluation

Here, we are interested in evaluating the scalability of DyRAM to a large number of receivers. Simulation scenarios performed on DyRAM and SRM are the same as those described in section 4.3 where simulations are run until a 200-packet file is correctly received by all the receivers. Remember that losses are generated during all the simulation time.

Figure 13 plots for different loss rates, the repair latency achieved by the two protocols as a function of the number of the subscribed receivers to the multicast session. We observe that DyRAM achieves a repair latency less than 1, thus less than the RTT to the source. This still being true even when increasing the number of the receivers which demonstrates the scalability properties of DyRAM. However, the repair latency in SRM is at least about twice the RTT to the source and increases with the number of the receivers to achieve ten times the RTT to the source when the number of the receivers exceeds 200.

Figure 14 plots for different loss rates, the repair implosion experienced by the two protocols when increasing the number of the receivers. We can observe that the repair implosion is always greater than 1 in SRM while it is less than 1 for DyRAM. More interesting for low losses (say 1% and 5%), the repair implosion decreases in DyRAM when the number of the receivers increases. This is due to the fact that more receivers are present per local group and chances to recover losses locally are bigger since loss rates are sufficiently low.


5 Conclusion

In this report, we presented a simulation-based evaluation of DyRAM where only the reliability mechanisms are considered. The associated congestion control protocol is dealt with in [5]. We tried to reflect the real characteristics of the Internet in the adopted network model. We proposed a loss model where both spatial and temporal correlation are considered and different loss scenarios are considered in the simulation experiences.


In this evaluation study, we evaluated the different active services proposed in DyRAM such as local recovery from the receivers and feedback suppression/aggregation. We found that DyRAM allows for less load in terms of consumed bandwidth on the different multicast nodes and that its different services allows for reducing the repair latency and implosion. The impact of varying the percentage of active routers in the multicast tree has also been considered. We showed that increasing the density of active routers allows for better performances. The comparative study conducted between DyRAM and SRM has also allowed for showing the load balancing feature of our approach. Finally, simulations with a large number of receivers have been performed in order to evaluate the scalability of DyRAM. Results have shown that DyRAM is scalable since it allows for a repair latency which is less than 1 even when increasing the number of the receivers.

References


- [1] Dah Ming Chiu, Stephen Hurst, Miriam Kadansky, and Joseph Wesley. Tram: A tree-based reliable multicast protocol. Technical Report TR-98-66, SUN, July 1998.
- [2] 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.
- [3] S. Kasera and S. Bhattacharya. Scalable fair reliable multicast using active services. *IEEE Network Magazine's Special Issue on Multicast*, 2000.




(a1)




(a2)




(b1)



(b2)



(c1)



(c2)

Figure 11: Load Balancing in DyRAM (left) and SRM (right), $p_l = 0.01$ (a) input, (b) output and (c) overall load.

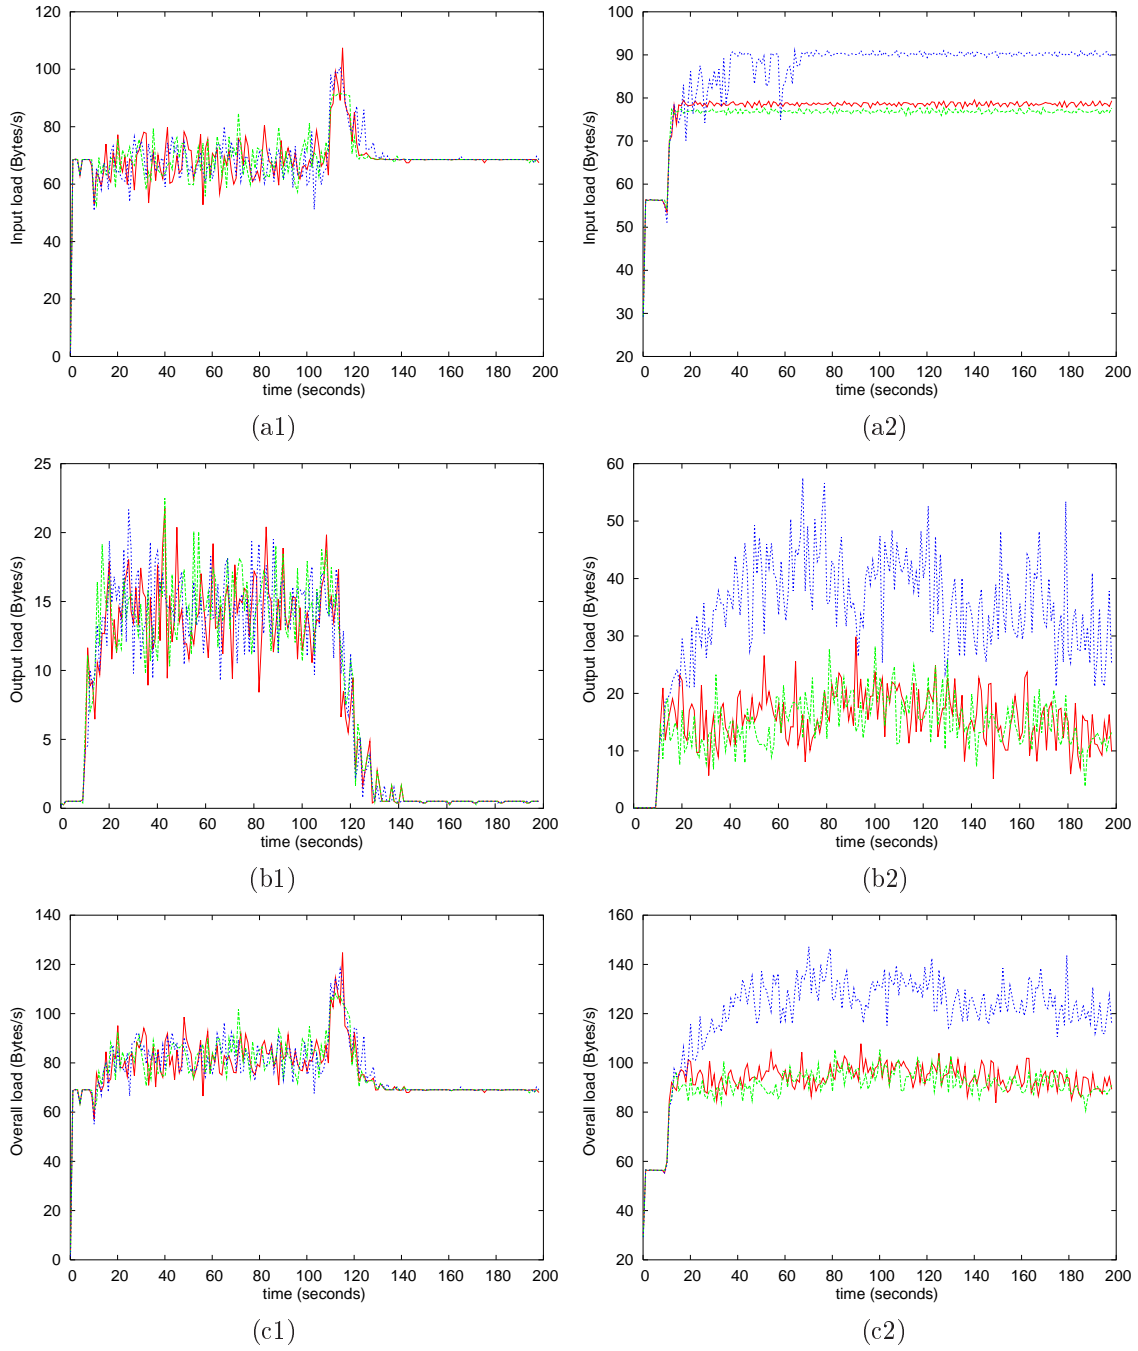


Figure 12: Load Balancing in DyRAM (left) and SRM (right), $p_l = 0.25$ (a) input, (b) output and (c) overall load.




Figure 13: Repair latency in SRM and DyRAM with losses at the tail links (a) $p_l = 0.01$, (b) $p_l = 0.05$, (c) $p_l = 0.1$ and (d) $p_l = 0.25$.

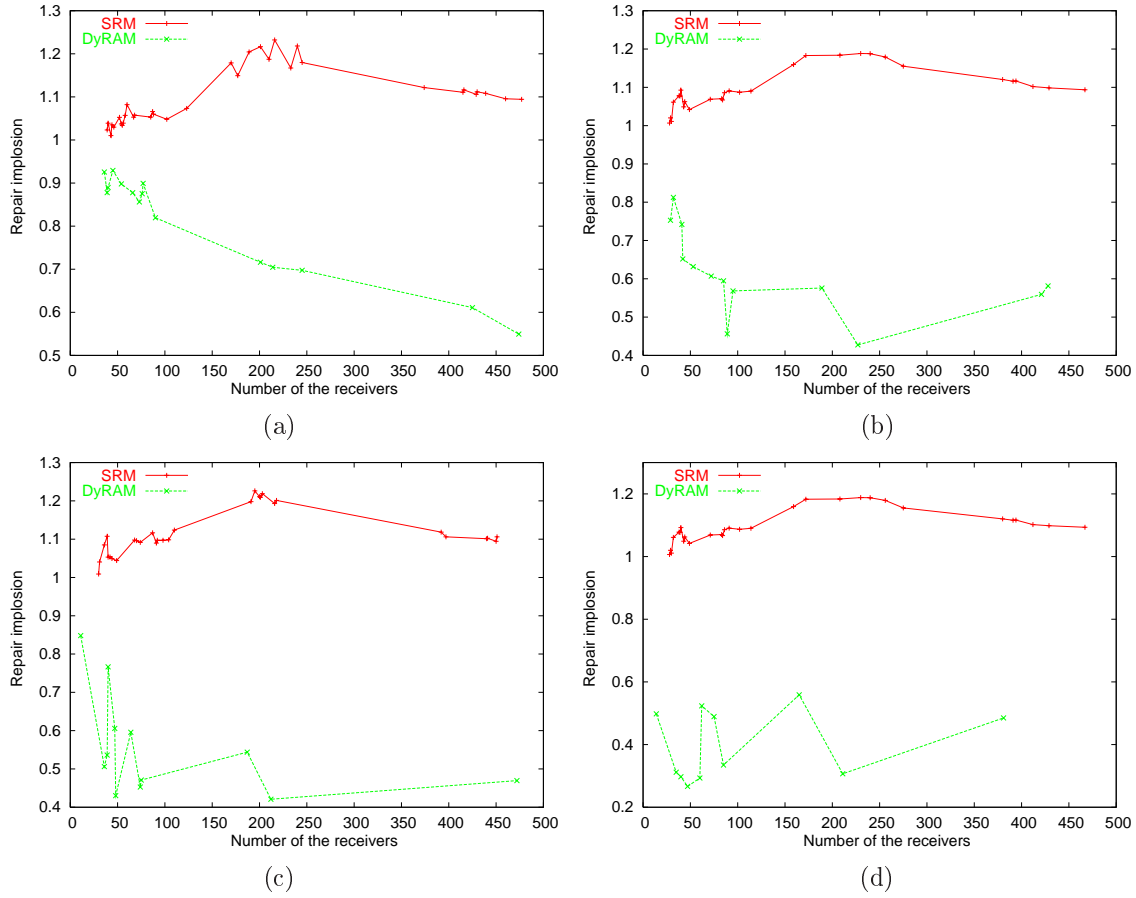


Figure 14: Repair implosion in SRM and DyRAM with losses at the tail links (a) $p_l = 0.01$, (b) $p_l = 0.05$, (c) $p_l = 0.1$ and (d) $p_l = 0.25$.

- [4] Moufida Maimour and Cong-Duc Pham. Dyrain : A reliable multicast protocol. Technical Report RR-4635, INRIA, January 2002. Also available as a LIP/ENS Research Report under 2003-05.
- [5] Moufida Maimour and Cong-Duc Pham. Amca : An active-based multicast congestion avoidance algorithm. Technical Report RR4689, INRIA, January 2003. Also available as a LIP/ENS Research Report under 2003-07.
- [6] Moufida Maimour and Cong-Duc Pham. Dyrain: an active reliable multicast framework for data distribution. *to appear in the Cluster Computing Journal*, 2004.
- [7] S. Paul and K. Sabnani. Reliable multicast transport protocol (RMTP). *IEEE JSAC, Spec. Issue on Network Support for Multipoint Communications*, 15(3), April 1997.
- [8] T. Speakman, J. Crowcroft, J. Gemmell, D. Farinacci, S. Lin, D. Leshchiner, M. Luby, T. Montgomery, L. Rizzo, A. Tweedy 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.
- [9] Timothy W. Strayer, Bert J. Dempsey, and Alfred C. Weaver. XTP – THE XPRESS TRANSFER PROTOCOL. *Addison-Wesley Publishing Company*, 1992.
- [10] Rajendra Yavatkar, James Griffioen, and Madhu Sudan. A reliable dissemination protocol for interactive collaborative applications. In *ACM Multimedia*, 1995.