



HAL
open science

Hybrid Checkpointing for Parallel Applications in Cluster Federations

Sébastien Monnet, Christine Morin, Ramamurthy Badrinath

► **To cite this version:**

Sébastien Monnet, Christine Morin, Ramamurthy Badrinath. Hybrid Checkpointing for Parallel Applications in Cluster Federations. 4th IEEE/ACM International Symposium on Cluster Computing and the Grid, Apr 2004, Chicago, IL, USA, United States. inria-00000991v1

HAL Id: inria-00000991

<https://inria.hal.science/inria-00000991v1>

Submitted on 10 Jan 2006 (v1), last revised 4 Mar 2016 (v3)

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.

Hybrid Checkpointing for Parallel Applications in Cluster Federations

Sébastien Monnet
IRISA
Sebastien.Monnet@irisa.fr

Christine Morin
IRISA/INRIA
Christine.Morin@irisa.fr

Ramamurthy Badrinath
Hewlett-Packard ISO, Bangalore, India¹
badrinar@india.hp.com

Abstract

Cluster federations are very useful for applications like large scale code coupling. Faults may appear very frequently, so we want to use checkpoints to be able to restart applications. To take into account the constraints introduced by clusters federation architecture, we propose a hierarchical checkpointing protocol. It uses synchronization inside clusters but only quasi-synchronous methods between clusters. Our protocol has been evaluate by simulation and fits well for applications that can be divided in modules with a lot of communications inside modules but few between them.

1 Introduction

We are working with cluster federations. These kinds of architectures are used when some parts of the code of a parallel application are running in different clusters. This can be due to security (a module of an application may need to run into its laboratory), to the hardware (a module may use sensors or specific hardware to display results), or simply for large scale needs. An example of a code coupling application running in a cluster federation could be different parallel simulation modules that sometimes need to communicate with each other.

There are lots of papers describing checkpoint / restart protocols inside a cluster in the literature. We want to take advantage of the high performance network (SAN) in the clusters and to take into account inter-cluster links which can be LANs or WANs. Considering the architecture, we want to use different mechanisms inside and between clusters. Our protocol is hierarchical : it uses coordinated checkpointing techniques inside clusters and communication-induced

¹This work was done when R. Badrinath was a visiting researcher at IRISA, on leave from IIT Kharagpur

checkpointing between them.

By simulating the protocol, we've seen that it works well with "code coupling" like applications.

The continuation of this document is made up as follows. The second Section shows the design principles, then Section 3 describes the hierarchical protocol combining coordinated and communication induced checkpointing (called HC³I checkpointing protocol thereafter). Section 4 shows a sample execution with the HC³I checkpointing protocol, Section 5 presents the evaluation. Before concluding, Section 6 is about related work.

2 Design Principles

This section presents the model considered in our work and the design principles of the HC³I checkpointing protocol.

2.1 Model

Application. We look after parallel applications like code coupling. Processes of these kinds of application can be divided into groups (modules). Processes inside a same group communicate a lot while communications between processes belonging to different groups are limited. Communications can be pipelined as in Figure 1 or they can consist of exchanges between two simulation modules for example.

Architecture. A cluster federation is a set of clusters linked by a WAN (Wide Area Network). Inter-cluster links can be dedicated ones or even Internet. It fits well with the kinds of application described above. Indeed each group of processes may run in a cluster where network links have small latencies and large bandwidths (System Area Network). We assume that a sent message will be received in

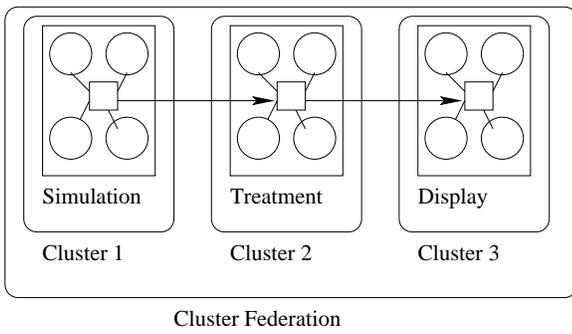


Figure 1. Application Model

an arbitrary but finite laps of time. This means that the network doesn't lose messages. This assumption implies that the fault tolerance mechanism should take care of messages in transit, they should not be lost.

Failure Assumptions. We assume that only one fault occurs at a time. However, the protocol can be extended to tolerate simultaneous faults as explain in the Future Work Section. The failure model is fail-stop. It means that when a node fails it won't send messages anymore. The protocol does not take into account neither omission nor byzantine faults.

2.2 Checkpointing Large Scale Applications in Cluster Federations

The basic principle of all checkpoint / rollback methods is to periodically store an application consistent state to be able to restart from there in case of failure in the future. A parallel application state is composed by the set of the states of all its processes. Consistent means that there is neither in transit message (sent but not received) nor ghost message (received but not sent) in the set of process states stored. A message generates a dependency. For example, Figure 2 presents the execution of two processes which both store their local state ($S1$ and $S2$). A message m is sent from process 2 to process 1. If the execution is restarted from the set of states $S1/S2$ the message m will have been received by process 1 but not sent by process 2 (ghost message). Process 2 will send m again which is not consistent because $S1$ happens before $S2$. The *happen before* relation is described in details in [4].

The last recorded consistent state is called the *recovery line*. [6] provides detailed information about the different checkpointing methods.

Inside a cluster we use a coordinated checkpointing method. This means there is a two phase commit protocol during which application messages are frozen. It ensures that the

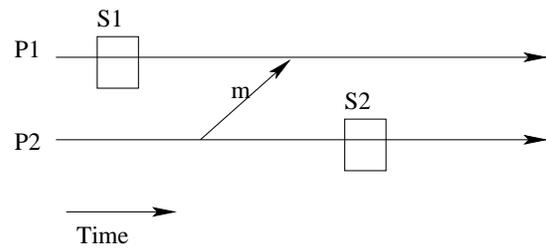


Figure 2. Dependency between two states

stored state (the cluster checkpoint) is consistent. Coordinated checkpointing is possible inside a cluster thanks to the high performance network (low latency and large bandwidth). It has already been done [8], [5],[11],[1] and is relatively easy to implement.

The large number of nodes and network performance between clusters don't allow a global synchronization. An independent checkpointing mechanism in which each cluster takes its Cluster Level Checkpoints (called *CLC* thereafter) won't fit. Tracking dependencies to compute the recovery line would be very hard at rollback time and clusters may rollback to very old *CLCs* (domino effect).

If we intend to log inter-cluster communications (to avoid dependencies), we need the piecewise deterministic (*PWD*) assumption. *PWD* assumption means that we are able to replay a parallel execution in a cluster that produces exactly the same messages as the first execution. This assumption is very strong. Replaying a parallel execution means detecting, logging and replaying all non-deterministic events. It is not always possible.

The assumption that inter-clusters communications are limited leads us to use a communication-induced method between clusters. This means each cluster takes *CLC* independently, but information is added to each inter-cluster communication. It may lead the receiver to take a *CLC* (called forced *CLC*) to ensure the recovery line progress. Communication-induced checkpointing seems to keep enough synchronization and can be efficient.

So our proposition is an hybrid protocol combining coordinated and communication-induced checkpointing (HC^3I).

3 Description of the HC^3I Checkpointing Protocol

In this section we first present the checkpointing mechanism inside a cluster. Then we describe mechanisms used to track inter-cluster dependencies and to decide when a *CLC* should be forced. Finally, we describe the rollback protocol and the garbage collector needed to eliminate *CLC* that are no longer useful.

3.1 Cluster Level Checkpointing

In each cluster, a traditional two phase commit protocol is used. An initiator node broadcasts (in its cluster) a *CLC request*. All the cluster nodes acknowledge the request, then the initiator node broadcasts a *commit*. Between the request and the commit messages, application messages are queued to prevent intra-cluster dependencies.

In order to be able to retrieve *CLC* data in case of a node failure, *CLCs* are recorded in the node own memory, and in the memory of one other node in the cluster. Because of this stable storage implementation, only one simultaneous fault in a cluster is tolerated.

Each *CLC* is numbered. Each node in a cluster maintains a sequence number (*SN*). *SN* is incremented each time a *CLC* is committed. This ensures that the sequence number is the same on all the nodes of a cluster (outside the two phase commit protocol). The *SN* is used for inter-cluster dependency tracking. Indeed, each cluster takes its *CLC* periodically, independently from the others.

3.2 Federation Level Checkpointing

If we look at our application model, communications between two processes in different clusters may appear. It implies dependencies between *CLCs* taken in different clusters. Dependencies need to be tracked to be able to restart the application from a consistent state.

Forcing a *CLC* in the receiver's cluster for each inter-cluster application message would work but the overhead would be huge as it would force useless checkpoints. In Figure 3, cluster 2 takes two forced *CLCs* (the filled ones) at message reception, and the application takes message into account only when the forced *CLC* is committed. *CLC2* is useful: in the event of a failure, a rollback to *CLC1/CLC2* will be consistent (*m1* would be sent and received again). On the other hand, forcing *CLC3* is useless: cluster 1 has not stored any *CLC* between its two message sending. In the event of a failure it will have to rollback to *CLC1* which will force cluster 2 to rollback to *CLC2*. *CLC3* would have been useful only if cluster 1 would have stored a *CLC* after sending *m1* and before sending *m2*.

Thus, a *CLC* is forced in the receiver's cluster only when a *CLC* has been stored in the sender's cluster since the last communication from the sender's cluster to the receiver's cluster. To do this, *CLCs* are numbered in each cluster with a *SN* (as described in previous section). The current cluster's sequence number is piggy-backed on each inter-cluster application message. To be able to decide if a *CLC* needs to be initiated, all the processes in each cluster need to keep the last received sequence number from each other cluster. All these sequence numbers are stored in a *DDV* (Direct Dependencies Vector, [2]).

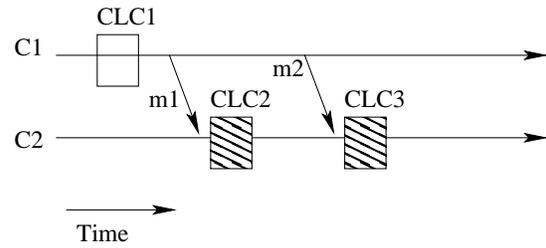


Figure 3. Limitation of the number of forced *CLCs*

$DDV_j[i]$ is the i^{th} *DDV* entry of cluster j , and SN_i is the sequence number of cluster i .

For a cluster j :

If $i=j$, $DDV_j[i]=SN_j$

If $i \neq j$, $DDV_j[i]=$ last received SN_i (0 if none).

Note that the size of the *DDV* is the number of clusters in the federation, not the number of nodes. In order to have the same *DDV* and *SN* on each node inside a cluster, we use the synchronization induced by the *CLC* two phase commit protocol to synchronize them. Each time the *DDV* is updated, a forced *CLC* is initiated which ensures that all the nodes in the cluster which takes a *CLC* will have the same *DDV* at commit time. The current *DDV* is stored with each *CLC*.

3.3 Logs to Avoid Huge Rollbacks

In order to limit the number of clusters that rollback, if the sender's cluster of a message doesn't fail while the cluster in which there's the receiver does, we don't want the sender's cluster to rollback. When a message is sent outside a cluster, the sender logs it optimistically in its volatile memory (logged messages are used only if the sender doesn't rollback). The message is acknowledged with the receiver's *SN* which is logged along with the message itself (inter-cluster messages are numbered). Next section explains which messages are replayed in the event of a failure.

3.4 Rollback

If a node fails inside a cluster, it is detected and the cluster rolls back to its last stored *CLC* (the description of the failure detector is out of the scope of this paper). One node in each other cluster in the federation receives a *rollback alert*. It contains the faulty cluster's *SN* that corresponds to the *CLC* to which it rolls back.

When a node receives such a *rollback alert* from another cluster with its new *SN*, it checks if its cluster needs to rollback by comparing its *DDV* entry corresponding to the faulty cluster to the received *SN*. If the former is greater or

equal than the latter its cluster needs to rollback to the first (the older) *CLC* which has its *DDV* entry corresponding to the faulty cluster greater or equal than the received *SN*. The node that has received the alert initiates the rollback.

If a cluster needs to rollback due to a received alert, it will send a *rollback alert* containing its new *SN* to alert all the other clusters. This is how the recovery line is computed.

Even if its cluster doesn't need to rollback a node receiving a *rollback alert* broadcasts it in its cluster. The nodes which have logged messages sent to a node in the faulty cluster acknowledged with a *SN* greater than the alert one or not acknowledged at all, re-send them.

Our communication induced mechanism implies that clusters need to keep multiple *CLC* and logged messages. They need to be garbage collected.

3.5 Garbage Collection

Our protocol needs to store multiple *CLCs* in each cluster in order to compute the recovery line at rollback time. The memory cost may become important. Periodically, or when a node memory saturates, a garbage collection is initiated. Our garbage collector is centralized. A node initiates a garbage collection, it asks one node in each cluster to send back its list of all the *DDVs* associated with the stored *CLCs*. Then it simulates a failure in each cluster and keep for each ones the worst *SN* to which they might rollback. It sends a vector containing all the worst *SNs* to one node in each cluster which broadcasts it in its cluster.

Each node removes the *CLCs* which have its cluster *DDV* entry smaller than the worst *SN* associated to its cluster. They also remove logged messages that are acknowledged with a *SN* smaller than the receiver's cluster worst *SN*.

4 Example

Figure 4 shows a sample execution on three clusters. It is composed of three parts. Each part is a snapshot of the execution. On each part, the execution time goes from left to right, each horizontal line represents a parallel execution on a cluster. The boxes are for the *CLCs*, the darker ones are forced *CLCs*. The corresponding *DDVs* are embedded in the *CLC*'s boxes.

The first part shows a normal execution until a failure appears in cluster 2. Notice that each cluster stores a first *CLC* which is the beginning of the application. Cluster 1 sends message *m1* to cluster 2, it sends its *SN* (1) along with *m1*. When receiving *m1*, cluster 2 compares the received *SN* with cluster 1 *DDV* entry (0). 1 is greater than 0, this forces cluster 2 to take a *CLC* before delivering *m1* to the application level. When receiving *m2* from cluster 1, cluster 2 doesn't have to initiate a *CLC*, the received *SN* (1) is equal to cluster 1 *DDV* entry in cluster 2. As for *m1*, we

see that *m3*, *m4* and *m5* force *CLCs* respectively on clusters 3,3 and 1. Notice that inter cluster messages are acknowledged with the local *SN* + 1 (the inter-cluster message will be delivered after the *CLC* is committed). Logged messages are not represented to keep the figure easy to read.

When a fault is detected in cluster 2, the whole cluster rolls back to its last stored *CLC*, its new *SN* is 3. So it sends a *rollback alert* with the *SN* 3 (second part of the figure). Cluster 1 doesn't have any cluster 2 *DDV* entry greater or equal to the received *SN* in its *DDVs* stored with the *CLCs*, it doesn't need to rollback. On the other hand, cluster 3 has to rollback to the first *CLC* that has its associated *DDV* containing cluster 2 entry greater or equal (equal in the sample) than the received *SN*. So cluster 3 has to rollback (third part of the figure) and sends an alert with its new *SN* (3).

Cluster 2 has never received messages from cluster 3 so its *DDVs* entries corresponding to cluster 3 are all equal to 0. It does not need to rollback. Cluster 1 has to rollback to its last *CLC* which has 4 as cluster 3 entry. It sends a rollback alert with its new *SN* (3) but no cluster has to rollback (due to the *DDV* lists).

Furthermore, the node in cluster 1 receiving the rollback alert from cluster 3, broadcasts it in its cluster. The node which has sent *m4* has to re-send it. It is acknowledged with *SN*=4, and the alert contains *SN*=3.

5 Evaluation

To evaluate the protocol, a discrete event simulator has been implemented. We have evaluated the overhead of the protocol in terms of network and storage cost first, then we observe what happens with different communication patterns. At last the garbage collector effectiveness and cost are evaluated.

5.1 Simulator

C++SIM library (<http://cxsim.ncl.ac.uk>) has been used to write the simulator. This library provides generic threads, a scheduler, random flows and classes for statistical analysis. Our simulator is configurable. The user has to provide 3 files: a *topology file*, an *application file* and a *timer file*. In the *topology file*, there is the number of clusters, the number of nodes in each cluster, the bandwidth and latency in each cluster and between clusters (represented as a triangular matrix) and the federation *MTBF* (Mean Time Between Failures). The *application file* contains, for each cluster, the nodes mean computation times, communication patterns between computations (represented by sends probabilities between nodes) and the application total time. At last, the *timers file* contains the delays for the protocol timers for each cluster (delays between two *CLCs*, garbage collection,...).

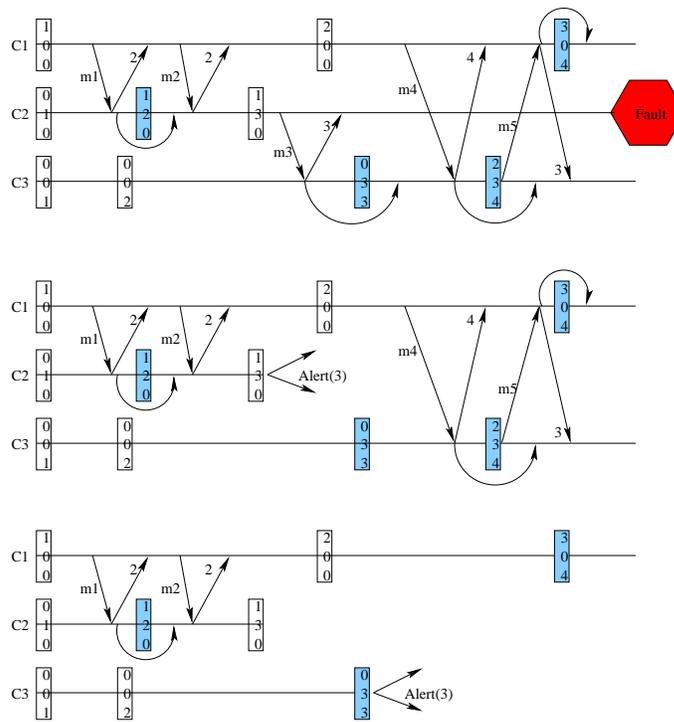


Figure 4. HC³I checkpointing protocol sample

The simulator is composed of four main threads. The thread *Nodes* takes the identity of all the nodes, one by one. The thread *Network* stores the messages and computes their arrival time. There is also a thread *Timers* to simulate all the different timers, and a thread *Controller* that controls the other threads (launches them, displays results at the end,...). Communication between threads is done by shared variables.

The simulator can be compiled with different trace levels. In the higher, we can observe each node action time-stamped (sends, receives, timer interruptions, log searches...). The lowest simulator output is statistical data, as messages count in clusters and between each cluster, number of stored *CLCs*, number of protocol messages,...

5.2 Network Traffic and Storage Cost

Evaluating network traffic and storage cost is very hard. It depends on how the protocol is tuned. If the frequency of unforced *CLCs* is low in a cluster, the *SNs* won't grow too fast so inter-cluster messages from this cluster would have a low probability to force *CLCs*. Reducing the protocol overhead becomes easy. If no *CLC* is initiated, the only protocol cost consists in logging optimistically in volatile memory inter-cluster messages and transmitting an integer (*SN*) with them. There is also a little overhead due to

message interception (between the network interface and the application).

To take advantage of the protocol, the timer that regulates the frequency of unforced *CLCs* in a cluster should be set to a value that is much smaller than the MTBF of this cluster. To illustrate this, the simulator simulates 2 clusters of 100 nodes. In both clusters the network is "Myrinet like" (10 μ s latency and 80MB/sec bandwidth). The clusters are linked by "Ethernet like" links (150 μ s latency and 100MB/sec bandwidth). The application total execution time is 10 hours. There are lots of communications inside each cluster and few between them. It can be a simulation running on cluster 0 and a trace processor on cluster 1 for example. The table below displays the number of messages (intra and inter-cluster).

Sender's Cluster	Receiver's Cluster	Message Count
Cluster 0	Cluster 0	2920
Cluster 1	Cluster 1	2497
Cluster 0	Cluster 1	145
Cluster 1	Cluster 0	11

Graph 5 and 6 show the number of forced and unforced committed *CLCs* in each cluster according to the delay between unforced *CLCs* in cluster 0 (x axis, in minutes). Cluster 1 delay between *CLCs* is set to infinite. Cluster 0 stores

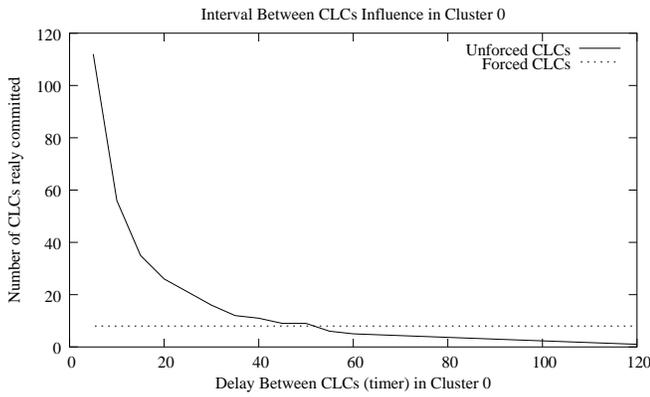


Figure 5. Number of CLCs in Cluster 0

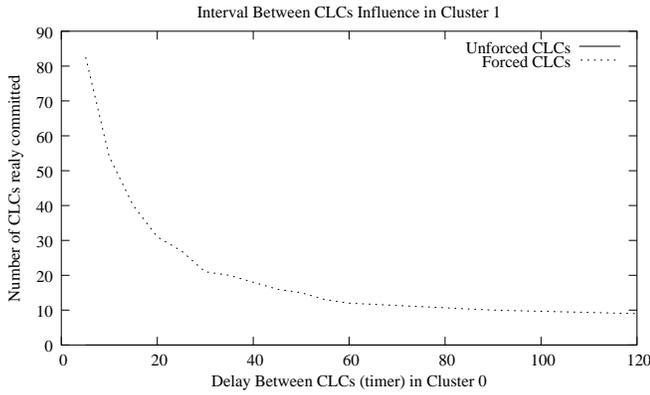


Figure 6. Number of CLCs in Cluster 1

some forced *CLCs* (8) because of the communications from cluster 1. This number of forced *CLCs* is constant - there are few messages from cluster 1. Notice that the total number of stored *CLCs* is smaller than $\frac{\text{total computation time}}{\text{delay between CLCs}} + \text{number of forced CLCs}$ because the timer is reseted when a forced *CLC* is established. Clusters store few more *CLCs*, but they are placed better (in time). Cluster 1 doesn't store any unforced *CLCs* as its timer is set to infinite, but it stores some forced *CLCs* induced by incoming communications from cluster 0. The number of these forced *CLCs* is proportional to the number of *CLCs* stored in cluster 0 - numerous messages come from cluster 0.

One may want to store more *CLCs* in cluster 1, if this cluster is intensively used and computation time is expensive for example. Graph 7 shows that cluster 0 (which "delay between *CLCs*" timer is set to 30 minutes) won't store more *CLCs* even if cluster 1 timer is set to 15 minutes. This is thanks to the low number of messages from cluster 1 to cluster 0.

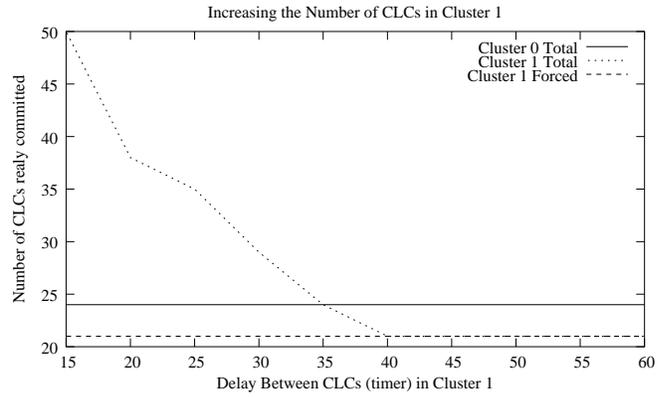


Figure 7. Impact of the Number of CLCs in Cluster 1

5.3 Communication Patterns

To better understand the influence of the communications patterns on the checkpointing protocol, Graph 8 shows what happens when the number of messages from cluster 1 to cluster 0 increases. Both cluster "delay between *CLCs*" timers are set to 30 minutes. The application is the same as in previous section except for the number of messages from cluster 1 to cluster 0 that is represented on the x axis.

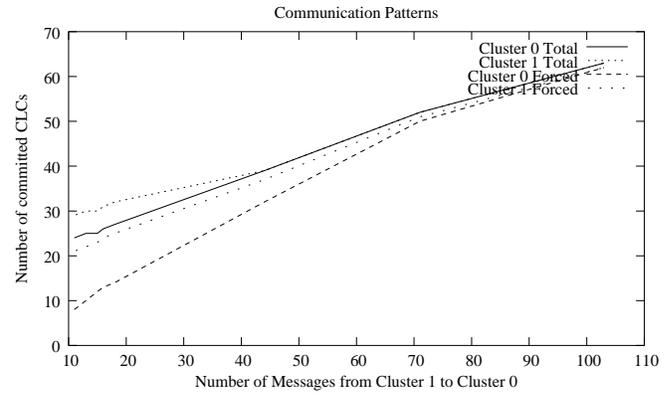


Figure 8. Increasing Communication from Cluster 1 to Cluster 0

The number of forced *CLCs* increases fast with the number of messages from cluster 1 to cluster 0. If the two clusters communicate a lot in both ways, *SNs* will grow very fast and most of the messages will induce a forced *CLC*.

5.4 Garbage Collection

A garbage collection has got a non negligible overhead. If N is the number of clusters in the federation, each garbage collection implies:

- $N-1$ inter-cluster requests
- $N-1$ inter-cluster responses which contains the list of all the *DDVs* associated to the stored *CLCs* in a cluster
- $N-1$ inter cluster collect requests
- A broadcast in each cluster

However, our hybrid checkpointing protocol may store multiple *CLCs* in each cluster. They can become very numerous. It also logs every inter-cluster application message. For the sample above, in the case of 103 messages sent from cluster 1 to cluster 0, without any garbage collection, there's 63 *CLCs* in each cluster. It means that each node in the federation stores 126 local states (its own 63 local states and the ones of one of its neighbor, because of the stable storage implementation).

If a garbage collection is launched every 2 hours, the maximum number of stored *CLCs* just after a garbage collection is 2 per cluster in this sample. Only oldest *CLCs* are removed, as explain in Section 3.5, so rollbacks won't be too big. The maximum number of logged messages during the execution in the sample above is 4 in both clusters.

The table below shows for each garbage collection the number of *CLCs* stored just before and just after the collection.

Cluster 0 Before	Cluster 0 After	Cluster 1 Before	Cluster 1 After
10	2	11	2
18	2	18	2
15	2	14	2
14	2	15	2

A second experimentation simulates an application that runs on three clusters. Clusters 0 and 1 have the same configuration as above. Cluster 2 is a clone of cluster 1. There's approximately 200 messages that leave and arrive in each cluster. The table below shows for each garbage collection the number of *CLCs* stored just before and just after the collection.

Cluster 0 (before)	30	48	54	38
Cluster 0 (after)	2	2	2	2
Cluster 1 (before)	50	80	78	64
Cluster 1 (after)	2	2	2	2
Cluster 2 (before)	50	80	78	64
Cluster 2 (after)	2	2	2	2

A tradeoff has to be found between the frequency of garbage collection and the number of *CLCs* stored.

6 Related Work

A lot of papers about checkpointing methods can be found in the literature. However, most of the previous works are related to clusters, or small scale architectures. A lot of systems are implemented at the application level, partitioning the application processes into steps. Our protocol is implemented at system level so that programmers don't need to write specific code. Moreover the protocol in this paper takes clusters federation architecture into account. This section presents several works that are close to ours.

Integrating Fault-Tolerance Techniques in Grid Applications. [9] does not present a protocol for fault tolerance but it describes a framework that provides hooks to help developers to incorporate fault tolerance algorithms. They have implemented different known fault tolerance algorithms and it seems to fit well with large scale. However, these algorithms are implemented at application level and are made for object-based grid applications.

MPICH-V. [3] describes a fault tolerant implementation of MPI. It is made for large scale architectures. All the communications are logged and can be replayed. This avoids all dependencies so that a faulty node will rollback, but not the others. But this means that strong assumptions upon determinism have to be taken. Our protocol doesn't need any assumption upon the application determinism, moreover it takes advantage of the fast network available in the clusters.

Hierarchical Coordinated Checkpointing Protocol. The work presented in [10] is the closest from ours. It proposes a coordinated checkpointing method, based on the two phase commit protocol. The synchronization between two clusters (linked by slower links) is relaxed. In [10], it's the coordinated checkpointing mechanism that is relaxed between clusters. It is not an hybrid protocol like ours. Our protocol is more relaxed, it is independent checkpointing if there are no inter-cluster messages.

7 Conclusion and Future Work

This paper describes an hybrid protocol combining coordinated and communication induced checkpointing methods. This approach is new and it works well with code-coupling like applications. It needs to be tuned according to the network and the application communication patterns. This protocol needs some improvements. Adding

some transitivity in the dependency tracking mechanism by sending the whole *DDV* instead of the *SN* should allow to take less forced checkpoints. This could allow to tolerate more communication patterns. The protocol should tolerate multiple faults in a cluster, this implies more redundancy in the stable storage implementation. It should tolerate simultaneous fault in different clusters, the garbage collector should take care of this. In the future we want to manage dynamism in order to be able to provide support for grid architectures. At last, the garbage collector should be more distributed.

A real implementation would be interesting to validate our protocol.

The algorithms and more details can be found in [7] (in French).

References

- [1] A. Agbaria and R. Friedman. Starfish: Fault-Tolerant Dynamic MPI Programs on Clusters of Workstations. *The Eighth IEEE International Symposium on High Performance Distributed Computing*, pages 167–176, August 1999.
- [2] R. Badrinath and C. Morin. Common mechanisms for supporting fault tolerance in DSM and message passing systems. Technical report, July 2003.
- [3] G. Bosilca, A. Bouteiller, F. Cappello, S. Djailali, G. Fedak, C. Germain, T. Herault, P. Lemarinier, O. Lodygensky, F. Magniette, V. Neri, and A. Selikhov. MPICH-V: Toward a Scalable Fault Tolerant MPI for Volatile Nodes. In *Proceedings of the IEEE/ACM SC2002 Conference*, pages 29–47, Baltimore, Maryland, November 2002.
- [4] K. Chandy and L. Lamport. Distributed Snapshots: Determining Global States of Distributed Systems. *ACM Trans. Computer Systems*, 3(1):63–75, February 1985.
- [5] M. Costa, P. Guedes, M. Sequeira, N. Neves, and M. Castro. Lightweight Logging for Lazy Release Consistent Distributed Shared Memory. In *Operating Systems Design and Implementation*, pages 59–73, October 1996.
- [6] M. Elnozahy, L. Alvisi, Y.-M. Wang, and D. Johnson. A Survey of Rollback-Recovery Protocols in Message-Passing Systems. *ACM Computing Surveys (CSUR)*, 34:375–408, September 2002.
- [7] S. Monnet. Conception et évaluation d’un protocole de reprise d’applications parallèles dans une fédération de grappes de calculateurs. Rapport de stage de DEA, IFSIC, Université de Rennes I, France, June 2003. In French.
- [8] C. Morin, A.-M. Kermarrec, M. Banâtre, and A. Geffaut. An Efficient and Scalable Approach for Implementing Fault Tolerant DSM Architectures. *IEEE Transactions on Computers*, 49(5):414–430, May 2000.
- [9] A. Nguyen-Tuong. *Integrating Fault-Tolerance Techniques in Grid Applications*. PhD thesis, Faculty of the School of Engineering and Applied Science at the University of Virginia, August 2000.
- [10] H. Paul, A. Gupta, and R. Badrinath. Hierarchical Coordinated Checkpointing Protocol. In *International Conference*

on Parallel and Distributed Computing Systems, pages 240–245, November 2002.

- [11] J. Rough and A. Goscinski. Exploiting Operating System Services to Efficiently Checkpoint Parallel Applications in GENESIS. *Proceedings of the 5th IEEE International Conference on Algorithms and Architectures for Parallel Processing*, October 2002.