

DISTRIBUTED OVERLOAD CONTROL FOR REAL-TIME REPLICATED DATABASE SYSTEMS

S. Saad-Bouzeffrane and C. Kaiser

CEDRIC laboratory, Conservatoire National des Arts et Métiers,

Tel: +33 (1) 40 27 25 83 - Fax: +33 (1) 40 27 22 96

292 rue Saint Martin, 75141 Paris Cédex 03, FRANCE

Email: {samia.bouzeffrane, kaiser}@cnam.fr

Key words: Firm real-time transactions, distributed database system, commit processing, processor overload, data replication.

Abstract: In order to meet their temporal constraints, current applications such as Web-based services and electronic commerce use the technique of data replication. To take the replication benefit, we need to develop concurrency control mechanisms with high performance even when the distributed system is overloaded. In this paper, we present a protocol that uses a new notion called *importance value* which is associated with each real-time transaction. Under conditions of overload, this value is used to select the most important transactions with respect to the application transactions in order to pursue their execution ; the other transactions are aborted. Our protocol RCCOS (*Replica Concurrency-Control for Overloaded Systems*) augments the protocol MIRROR, a concurrency control protocol designed for firm-deadline applications operating on replicated real-time databases in order to manage efficiently transactions when the distributed system is overloaded. A platform has been developed to measure the number of transactions that meet their deadlines when the processor load of each site is controlled.

1 Introduction

Current applications, such as Web-based services, electronic commerce, mobile telecommunication system, etc., are distributed in nature and manipulate time-critical databases. In order to enhance the performance and the availability of such applications, one of the main techniques is to replicate data on multiple sites of the network. Therefore, the major issue is to develop efficient replica concurrency control protocols that are able to tolerate the overload of the distributed system. In fact, if the system is not designed to handle overloads, the effects can be catastrophic and some primordial transactions of the application can miss their deadlines. While many efforts have been made in the management of transactions for replicated databases in the real-time context (Son, 1987), (Son and Kouloumbis, 1993), (Son and Zhang, 1995), (Ulusoy, 1994), (M. Xiong and Stankovic, 2002), no work deals with protocols that manage distributed real-time databases and simultaneously control the overload of the system. Some researches have dealt with the scheduling of tasks under overload conditions when the real-time system is centralized such as in (Delacroix, 1996) and (Koren and

Shasha, 1995), or distributed as in (Kaiser and Santellani, 1998). In (Delacroix and Ménival, 2000), a protocol that controls the processor overload has been integrated in RT-Linux.

In this paper, we focus on the design of *one-copy serializable* concurrency control protocols for replicated real-time databases when the distributed system is overloaded. The overload occurs when the computation time of transactions set exceeds the time available on the processor and then the deadlines can be missed. Our study is concerned by "firm-deadline" transactions because many current applications such as Web-based services use communication protocols with timeout features. In firm-deadline applications, each transaction that misses its deadline is useless and is then aborted immediately.

We present in this paper, a replica concurrency-control protocol that manages the overload on each site of the distributed system by choosing only the "most important" transactions with respect to the application transactions, to maintain their execution ; the other transactions are simply aborted. An *importance value* is defined for each transaction in order to distinguish the most important transactions for which it is essential to maintain the execution from

the rest of transactions that may be aborted. Our aim is to alleviate the processor load. The importance criterion is used only to control the processor load since that ready transactions of each site are scheduled using a local scheduling algorithm such as EDF algorithm (Liu and Leyland, 1973). Our protocol RCCOS (*Replica Concurrency-Control for Overloaded Systems*) augments MIRROR (*Managing Isolation in Replicated Real-time Object Repositories*) (M. Xiong and Stankovic, 2002) a concurrency control protocol designed for replicated real-time databases that adds data conflict-resolution mechanisms to the optimistic two-phase locking (O2PL) (Carey and Livny, 1991) designed for a non real-time context. An experimental platform is being developed to implement RCCOS and MIRROR protocols in order to compare their performance under conditions of overload.

The remainder of this paper is organized as follows : Section 2 presents the related work. In Section 3, is described the model used whereas in Section 4 are described RCCOS protocol and MIRROR protocol on which it is based. The protocol implementation is described in Section 5. Some elements of proof are given in Section 6. Before concluding the paper in Section 8, we present, in Section 7, some performance results measured on an experimental platform that has been developed for this aim.

2 Related work

The protocols designed to manage distributed transactions in replicated databases (Anderson et al., 1998), (Agrawal et al.,), (Bernstein et al., 1987), (Kemmer and Alonso, 1998), (Pedone et al., 1998) are not suitable to real-time context because they are prone either to blocking or to message overhead or to both. This is incompatible with the respect of real-time constraints. Authors (eg., (Son, 1987), (Son and Kouloumbis, 1993), (Son and Zhang, 1995), (Ulusoy, 1994)) have then begun to work on concurrency control for replicated Distributed Real-Time Data Base Systems (DRTDBS) and have designed protocols that take into account the transactions real-time constraints. For example, a multiversion technique is proposed in (Son, 1987) to increase the degree of concurrency. In (Son and Kouloumbis, 1993), (Son and Zhang, 1995), a real-time scheduling is integrated in replication control. However, there is no study on how to control concurrency of transactions when the system is overloaded and to reach the double objective of any DRTDBS, which is : to maintain the database consistency while allowing to each transaction to meet its deadline. While no work has dealt with the management of real-time transactions in overloaded systems, some authors have designed

real-time scheduling algorithms that are resistant to the effects of system overload (Cottet et al., pear). In fact, some algorithms deal with periodic task sets and allow the system to handle variable computations times which cannot always be bounded (Atlas and Bestavros, 1998), (Mok and Chen, 1997), (Buttazo et al., 1998). Other algorithms deal with hybrid task sets where tasks are characterized with an importance value (Delacroix, 1996), (Koren and Shasha, 1995), (Kaiser and Santellani, 1998). All these policies deal with overloads to provide deadline missing tolerance. In (Baruah and Haritsa, 1997), Baruah et al. propose ROBUST an on-line uniprocessor scheduling algorithm that performs for a large range of slack factors during overload. Slack factor is defined as the ratio between the task deadline and its execution time.

Among the techniques that use the concept of importance, Kaiser et al., in (Kaiser and Santellani, 1998), have proposed a protocol to control the tasks load in a distributed real-time system. Tasks are assigned values used to define the importance degree of each task with respect to the application-tasks set. In order to decrease the tasks load, only the tasks declared "important" by the protocol have their execution maintained, the other tasks considered as less important are aborted. Delacroix et al., in (Delacroix and Ménéval, 2000), have implemented an overload controller that has been integrated to RT-Linux. In their task model proposed, each task is made up of several execution modes : the *normal mode* which is a mode executed when the task begins to execute. It takes care of the normal execution of the task. And the *survival modes* which are executed when the task is cancelled by the overload resorption or when it misses its deadline.

3 The model

In our model, we consider only firm real-time transactions. In other words, only transactions that fully complete execution before their deadlines are considered to be successful, whereas transactions that miss their deadlines are considered worthless and are immediately discarded without being executed to completion (Haritsa et al., 1991).

The distributed system is composed of :

- a master site on which each global transaction is submitted and
- sites that have replicas of parts of a distributed database.

Each submitted transaction T_i is decomposed into n subtransactions ST_{ij} that are sent to sites called *participant sites*. All these subtransactions are mandatory, that is, T_i commits only if all its subtrans-

actions commit locally within their sites. On each participant site, subtransactions are managed by a *cohort process*. The participant site of a subtransaction ST_{ij} is the one that manipulates the least important transactions among the sites that have replicas of data-items needed by ST_{ij} . The sites that are not participant sites but that have replicas of the database are called *updaters* (see Figure 1). All the sites, including the master, are partially connected, but the graph representing the network is a connected graph, that is, from any site it is possible to access another site directly or indirectly. A global transaction is characterized by its arriving time and its deadline while, within a participant site, a subtransaction is characterized by its arriving time, its execution duration and its deadline. The execution time must be lower than the relative deadline. Moreover, the subtransaction deadline is identical to the deadline of the global transaction to which the subtransaction belongs.

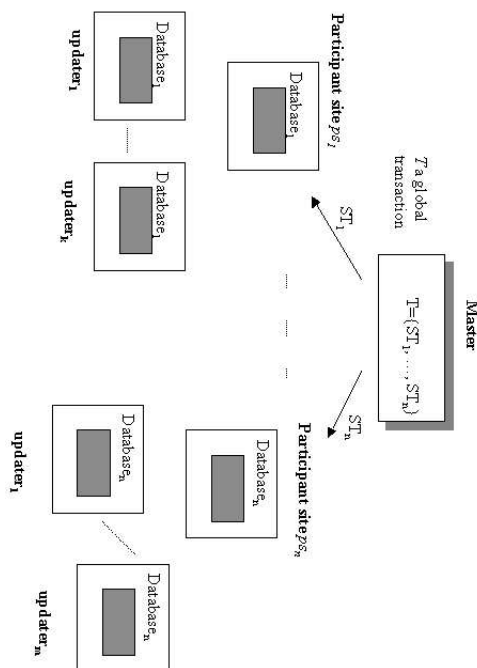


Figure 1: A replicated distributed database system

4 Protocol description

Since our protocol RCCOS augments the MIRROR protocol in order to handle the distributed-system overload, we first review the principle of MIRROR before describing the proposed protocol.

4.1 MIRROR protocol

The MIRROR protocol described in (M. Xiong and Stankovic, 2002) extends the implementation strategy for centralized optimistic concurrency-control algorithms proposed in (Huang et al., 1991) to handle data distribution and replication.

Transactions follow three steps during their execution : read, validation and write. In the read phase, cohorts acquire locks to access data items on their local sites. The updating of replicas is deferred to the end of transaction, that is, to the commit process.

In the validation phase, a cohort that receives a PREPARE message from its master begins a local validation. If it fails, it sends an ABORT message to the master. Otherwise it sends PREPARE messages containing relevant updates to its updater sites. Each updater site that receives a PREPARE message requests write locks to update the data in its local work-area. As soon as the updates have been performed, a PREPARED message is returned to the cohort. The cohort sends a PREPARED message to the master after it receives PREPARED messages from all its updaters. As soon as the master receives PREPARED messages from all its cohorts, it validates globally the transaction by sending COMMIT messages to all the cohorts.

The write phase begins when a cohort receives a COMMIT message. After a cohort writes on the database, it sends COMMIT messages to its updaters to update the database replicas.

To resolve data conflict, two mechanisms are proposed by MIRROR based on the following observation : "it is very expensive to abort a transaction when it is near to completion because all the resources consumed by the transaction are wasted" (M. Xiong and Stankovic, 2002).

- **Priority Blocking (PB)** mechanism : any transaction that cannot acquire a lock is inserted in a *lock request queue* until the required lock is released. The queue is ordered by transaction priority.
- **Priority Abort (PA)** mechanism : if a transaction T has higher priority than R and T requests a lock already held by R then instead of aborting R in order to execute T as it is done usually in real-time systems ; R will resume its execution if R is executing in its later stages. In this case, T will not wait too long since R is ending its execution. On the other hand, if R is executing in its first stages then it is aborted in order to execute T .

The rule introduced by Xiong et al. in (M. Xiong and Stankovic, 2002) allows to use PA if the lock holder has not reached yet a *demarcation point* and to use PB after the demarcation point is reached. The demarcation point corresponds :

- to the reception from the master of a PREPARE

message for a subtransaction ST_{ij} , if the receiver is a participant site and

- to the acquisition of all the local write locks by the local subtransaction if the site is a replica updater.

4.2 RCCOS Protocol

Our protocol is similar to MIRROR protocol under normal (non-overload) conditions. But when the system is overloaded because on each site, the computation time of transactions set exceeds the time available on the processor, we have to favour the executions of the most important transactions according to the application-transactions set. The favoured transactions are those which undergo less timing faults and which are less dropped. It is then necessary to have a means to designate the most important transactions. This means is presented in the form of a parameter called *importance value*, as described in (Kaiser and Santellani, 1998) for distributed real-time systems.

4.2.1 Transaction Importance

We associate, to each global transaction, a positive integer that represents the importance value of the transaction according to the application-transactions set. The importance value is intrinsic to the application, so it can be fixed by the application developer. Each transaction submitted to the master is characterized by a deadline which defines its urgency and by a value which defines the importance of its execution, with respect to the other transactions of the real-time application. The importance (or criticality) of a transaction is not related to its deadline; thus two different transactions which have the same deadline may have different importance values. We define the importance value of a ready-transactions queue as the importance value of the most important transaction of the queue. If $readyQueue_s$ is the ready-transactions queue of a site s , then the importance value of the site denoted Imp_s is : $Imp_s = Imp(readyQueue_s) = \text{Max} \{Imp(T_i)\}; \forall T_i \in readyQueue_s$.

4.2.2 RCCOS Protocol Principle

RCCOS protocol is applied to the distributed DBMS (Data Base Management System) model described in section 3 It is upstream of MIRROR protocol. When the master receives a transaction T_i to execute, it splits it into subtransactions. For each subtransaction, it tries to find a participant site. Since the database is replicated on many sites, in normal (non-overload) situation, for each subtransaction ST_{ij} any site among those that have replicas of the data items needed by ST_{ij} can be chosen to designate the participant site. But when the system is overloaded, the first step is

to use a placement policy based on the importance concept. In other words, before executing subtransactions we have to find for each subtransaction ST_{ij} the participant site that manipulates the least important transactions, among $m_{ST_{ij}}$ sites that are candidates to be the cohorts. Particularly, if the importance values of the $m_{ST_{ij}}$ sites, collected by the master, are equal to the importance value of the global transaction T_i then T_i is aborted by the master because the importance values must be distinct within a site so as to distinguish between transactions with high importance values from transactions with low importance values. The site that has the least important transactions is chosen to accept the execution of ST_{ij} , hence to become the cohort, while the remaining sites that have replicas of the database will be the updaters. Our choice is motivated by the fact that, in overload conditions, we have to carry out the most important transactions in the system : it is then better to submit a subtransaction to a site that handles transactions with low importance values than the reverse.

In order to find a possible cohort for a subtransaction ST_{ij} , the master broadcasts an INITIATE message as well as the global transaction T_i to all the sites that have replicas of the data items needed by ST_{ij} . When a site receives an INITIATE message, it sends its importance value. As soon as the master has collected the importance values of all the concerned sites, it declares as a cohort the site that has the lowest importance value by sending to it a COMPUTE message in order to execute ST_{ij} .

The site that receives a COMPUTE message is considered as the cohort. The other sites that fail to be the cohorts will be the updaters. The cohort, in this case, has to stabilize the new ready-transactions queue. The ready queue contains current ready transactions as well as the subtransaction ST_{ij} added to the queue during the treatment of the COMPUTE message. The stabilization process consists in, when the system is overloaded, privileging the transactions that have high importance values. The transactions that have low importance values are released from the queue and aborted one after the other until the processor laxity is a positive value. The processor laxity, at time t , is the maximum time the processor may remain idle, after t , without causing a transaction to miss its deadline.

In order to stabilize the ready queue, the transaction that will be aborted is chosen among transactions with low importance values while applying PA real-time mechanism introduced by MIRROR protocol. Hence, the overload is resorbed by a rejection policy based on removing transactions that have not reached their demarcation point and that have minimal importance values. A site considered as the cohort for a subtransaction can be, at the same time, the updater for another subtransaction. Therefore, the demarcation

point for a subtransaction may correspond to the reception of PREPARE messages or to the acquisition of all the required write locks depending on whether the site is viewed as the cohort or as the updater for this subtransaction.

5 RCCOS Algorithm

The algorithm uses some variables and data structures. We describe them in the following subsections.

5.1 Notations

- T_i : denotes a global transaction submitted to the master site. It is composed of k subtransactions denoted by $\{ST_{i1}, ST_{i2}, \dots, ST_{ik}\}$ that will be executed on k participant sites.
- Imp_{T_i} : a positive integer value that denotes the importance value of T_i . The importance values of the subtransactions are equal to the importance value of the global transaction to which they belong. Hence, $Imp_{ST_{ij}} = Imp_{T_i} \forall j \in \{1, k\}$.
- Imp_s : is the importance value of a site s . It corresponds to the highest importance value of the ready transactions within s .
- $Dead_{T_i}$: denotes the absolute deadline of T_i . $Dead_{T_i}$ is also the deadline of the subtransactions of T_i .
- $E_{ST_{ij}}$: denotes a set of sites that have replicated data items needed by ST_{ij} . This set includes the cohort and the updaters.
- $MinImp_{ST_{ij}}$: is the lowest importance value of $E_{ST_{ij}}$ and s_{min} is the site that has $MinImp_{ST_{ij}}$ as the importance value.
- $readyQueue$: is a list that contains ready transactions. Each element of the queue has the structure of $readyTrans$ as described below. The transactions are sorted by increasing their absolute deadlines. That is, the highest-priority transaction has the nearest deadline.
- $readyTrans$: is the structure of each element of the $readyQueue$ ($subtrans$: the subtransaction concerned, $globalTrans$: the global transaction to which belongs $subtrans$, $absDead$: the deadline of $subtrans$, $remExecTime$: initialized to the total execution time of $subtrans$ and is decreased during the execution to store the pending execution time of $subtrans$, $next$: the next element in the queue).
- $importanceQueue$: contains ready transactions that are sorted by increasing their importance values, that is, the first transaction of the queue is the least important one. Each element of the queue has the structure of $impTrans$ as described below.
- $impTrans$: is the structure of each element of the $importanceQueue$ ($subtrans$: the subtransaction

- concerned, $globalTrans$: the global transaction to which belongs $subtrans$, Imp : the importance value of $subtrans$, $next$: the next element in the queue).
- $currentReady$: is the current ready transaction dealt with when reading $readyQueue$ elements.
- $computTime$: is an integer value that cumulates the pending execution time of treated transactions.
- $deleteTime$: is an integer value that represents the pending execution time of the transactions that will be aborted in order to alleviate the processor load.
- $nbResp$: an integer value that counts the number of messages not received yet.

5.2 Exchanged messages

- INITIATE(ST_{ij}, T_i) : is a message sent by the master to find the site that manipulates the least important transactions, in order to be the cohort.
- IMP(ST_{ij}, T_i, Imp_s) : is a message sent by a site s to inform the master about its current importance value.
- COMPUTE($ST_{ij}, T_i, Imp_{T_i}, Dead_{T_i}$) : this message is sent by the master to the cohort in order to trigger the execution of ST_{ij} on the cohort.
- ABORT(ST_{ij}, T_i) : if the stabilization mechanism chooses ST_{ij} as a victim, an ABORT message is sent to the master that has to abort T_i after being sent ABORT messages to the other cohorts. The cohorts, of course, propagate this message to the updaters if the updates have been prepared. We note that no message ABORT is sent by the cohort if ST_{ij} misses its deadline because the deadline expiration is detected locally on the master, on the cohorts and on the updaters.

5.3 Text of the Algorithm

In order to designate a cohort for a subtransaction ST_{ij} , i.e., a site that has the lowest importance value in order to execute ST_{ij} of a global transaction T_i , the master broadcasts INITIATE messages to all the sites that have replicas of data items needed by ST_{ij} . Figure 2 summarizes the steps of the algorithm by a diagram and the details of the algorithm are given below. We first describe the behaviour of the master then that of the cohort.

Within the master :

```

WHEN a transaction  $T_i$  composed of  $k$  subtransactions is submitted to the master :
DO
 $\forall j \in \{1, k\}$  :
  BEGIN
     $nbResp_{ST_{ij}} = \text{Cardinal}(E_{ST_{ij}})$ ;
     $MinImp_{ST_{ij}} = +\infty$ ;
  
```

```

     $\forall s \in E_{ST_{ij}}$  :
        send INITIATE( $ST_{ij}$ ,  $T_i$ ); /* broadcasts
        INITIATE messages to all the replicas sites */
    END
END DO

```

WHEN the master receives IMP(ST_{ij} , T_i , Imp_s)
from a site s :

```

DO
     $nbResp_{ST_{ij}} = nbResp_{ST_{ij}} - 1$ ;
    1 IF ( $Imp_s \neq Imp_{T_i}$ ) AND ( $Imp_s < MinImp_{ST_{ij}}$ )
        THEN
             $MinImp_{ST_{ij}} = Imp_s$ ;
             $s_{min} = s$ ;
        2 ENDIF
    IF ( $nbResp_{ST_{ij}} = 0$ )
        THEN
            /* all the IMP messages have been received */
            send COMPUTE( $ST_{ij}$ ,  $T_i$ ,  $Dead_{T_i}$ ,  $Imp_{T_i}$ )
            to  $s_{min}$ ;
            /*  $s_{min}$  is the cohort of  $ST_{ij}$  */
        ENDIF
    END DO

```

Within any site :

WHEN a site s receives INITIATE(ST_{ij} , T_i) from
the master :

```

DO
    /* returns its importance value */
     $Imp_s = \text{Max} \{ Imp(ST_k) \} \forall ST_k \in readyQueue_s$ ;
    send IMP( $ST_{ij}$ ,  $T_i$ ,  $Imp_s$ ) to the master;
END DO

```

Within the cohort :

WHEN the cohort receives COMPUTE(ST_{ij} , T_i ,
 $Dead_{T_i}$, Imp_{T_i}) from the master :

```

DO
    /* stabilizes the local ready queue */
     $t = \text{current-time}()$ ;
    insert ( $ST_{ij}, T_i, Dead_{T_i}$ ) in  $readyQueue$ ;
    insert ( $ST_{ij}, T_i, Imp_{T_i}$ ) in  $importanceQueue$ ;
     $ov = \text{OVERLOAD-PROCEDURE}(t)$ ;
    /*  $ov$  stores the overload value returned */
    IF ( $ov \neq 0$ )
        THEN
             $deleteTime = 0$ ;
            3  $trans = \text{first transaction of } importanceQueue$ ;
            WHILE ( $trans \neq \text{null}$ )
                DO
                    4 IF ( $trans$  has not reached its demarcation
                    point) THEN
                        release  $trans$  from  $importanceQueue$ ;
                        release  $trans$  from  $readyQueue$ ;
                        abort  $trans$ ;
                        send ABORT( $trans.subtrans, trans.globalTrans$ )

```

to the master ;

```

         $deleteTime = deleteTime + trans.remExecTime$ ;
        IF ( $deleteTime \geq ov$ ) THEN exit;
        ENDIF
    ENDIF
     $trans = trans.next$ ;
    5 END WHILE
    ENDIF
END DO

```

OVERLOAD-PROCEDURE (integer t)

```

BEGIN
    int  $loadValue = 0$ ;
    int  $computTime = 0$ ;
     $currentReady = \text{first transaction of } readyQueue$ ;
    WHILE ( $currentReady \neq \text{null}$ )
        DO
            /* if the laxity condition of  $currentReady$  is a
            negative value */
            IF ( $computTime +$ 
                 $currentReady.remExecTime$ 
                 $currentReady.absDead - t$ )
                THEN /* memorize the overload value */
                     $loadValue = loadValue +$ 
                     $currentReady.remExecTime$ ;
                ELSE  $computTime = computTime +$ 
                     $currentReady.remExecTime$ ;
                ENDIF
             $currentReady = currentReady.next$ ;
        END WHILE
    return( $loadValue$ );
END

```

6 Experimental platform

As it is shown in Figure 3, transactions are submitted via Web pages to the master which trigger sub-transactions on different replica sites while handling the sites overload. On each site, there are three modules :

- the *scheduler* module that uses EDF algorithm to schedule local subtransactions,
- the *replicas data manager* that resolves data conflict and that handles replicated data and
- the *overload controller* module that implements RCCOS protocol.

This platform has been developed in Java and JDBC/ORACLE. The first experiments (see Figure 4) that have been done show that among transactions that meet their deadlines, the majority concerns transactions with high importance values. This behaviour can be explained as follows : the great percentage of deadline meeting is due to transactions with high importance values because RCCOS favours this kind

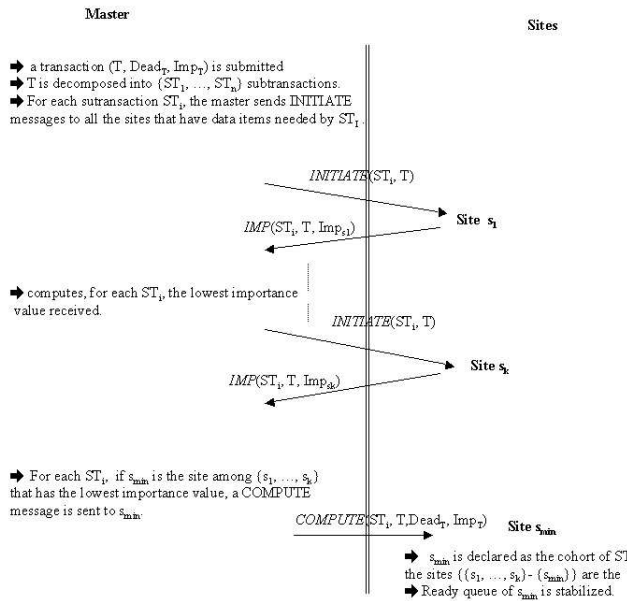


Figure 2: The steps followed to place a subtransaction on a cohort

of transactions. The little percentage of transactions with low importance values that meet their deadlines is due to the transactions that have reached their demarcation point, hence they are not aborted when applying stabilization process.

7 Conclusion and future work

In this paper, we have addressed the problem of managing firm-deadlines transactions that access replicated data in a distributed system that may be overloaded. For this environment in which many current time-critical applications operate, specially Web-based services, we proposed a novel protocol called RCCOS that can be easily integrated in current systems to handle overload processors without altering the database consistency which is the main objective of DRDDBSs. The main idea of the protocol is to associate an importance value to each submitted transaction in order to favour, when the system is overloaded, the executions of the most important transactions according to the application-transactions set. In overload conditions, each subtransaction of the global transaction is executed on a site that has the lowest importance value among those that have replicas of the data items needed by the subtransaction. Moreover, the participant site has to stabilize its ready queue, that is, it has to manage the processor overload by maintaining the execution of the most important transactions and by aborting the

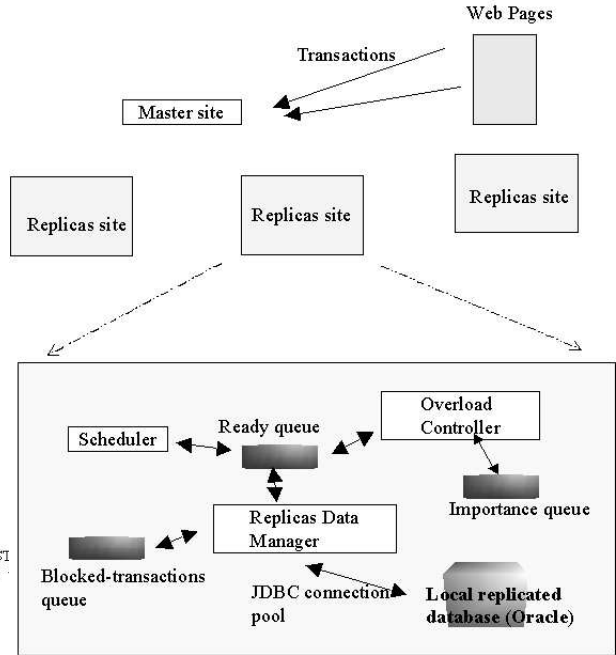


Figure 3: The different modules of the experimental platform

others. Priority Abort real-time mechanism of MIRROR protocol has been added to RCCOS protocol when proceeding to stabilization so as to favour the deadline meeting for the transactions that have high importance values and/or that are near to completion. The first experiments that have been done on our platform show that the majority of transactions that meet their deadlines have high importance values. We are currently working on integrating MIRROR protocol into our simulation platform to investigate the performance improvements that may arise from RCCOS protocol when comparing the performance of MIRROR and RCCOS protocols according to the same processor workload. We will also investigate mechanisms to minimize communications and response times between the cohort and the updaters when managing real-time database replicas.

Acknowledgement

Special thanks to SAAD Abdelkrim who helps us in the development of the experimental platform.

REFERENCES

- Agrawal, D., Abbadi, A. E., and Steinke, R. C. Epidemic algorithms in replicated databases (extended abstract). In *Proc. of the Sixteenth ACM SIGACT-SIGMOD-*

SIGART Symposium on Principles of Databases Systems.

- Anderson, T., Breitbart, Y., Korth, H. F., and Wool, A. (1998). Replication, consistency and practically: Are these mutually exclusive? In *Proc. of the ACM SIGMOD International Conference on Management of Data*, pages 484–495, Seattle, WA, USA.
- Atlas, A. and Bestavros, A. (1998). Statistical rate monotonic scheduling. In *IEEE Real-Time Systems Symposium*.
- Baruah, S. K. and Haritsa, J. R. (1997). Scheduling for overload in real-time systems. In *IEEE Trans. on Computers*, volume 46(9), pages 1034–1039.
- Bernstein, P., Hadzilacos, V., and Goodman, N. (1987). Concurrency control and recovery in database systems. In *Addison Wesley*, Massachussets.
- Buttazo., G. C., Lipari, G., and Abeni, L. (1998). Elastic task model for adaptive rate control. In *IEEE Real-Time Systems Symposium*, Madrid.
- Carey., M. and Livny, M. (1991). Conflict detection trade-offs for replicated data. In *ACM Transactions on Database Systems*, volume 16, pages 703–746.
- Cottet, F., Delacroix, J., Kaiser, C., and Mammeri, Z. (to appear). *Scheduling in Real-Time Systems*. Edition J. Wiley and Son.
- Delacroix, J. (1996). Towards a stable earliest deadline scheduling algorithm. In *Real-Time Systems Journal*, volume 10(3), pages 236–291.
- Delacroix, J. and M'énival, C. (2000). Intégration d'un contrôle de charge par importance au sein du système rt-linux. In *RTS'2000 Conference*, pages 47–63, Paris.
- Haritsa, J., Carey, M. J., and Livny, M. (1991). Earliest deadline scheduling for real-time database systems. In *1991 IEEE Real-Time Systems Symp.*
- Huang, H., Stankovic, J. A., Ramamritham, K., Towsley, D., and Purimetla, B. (1991). Experimental evaluation of real-time optimistic concurrency control schemes. In *the 17th International Conference on Very Large Data Bases*, Barcelona.
- Kaiser, C. and Santellani, C. (1998). P'etrarque. une plateforme d'expérimentation pour l'ordonnement temps réel strict d'applications réparties. In *Technique et Science Informatique Journal*, volume 17(1), pages 39–62.
- Kemme, B. and Alonso, G. (1998). A suite of database replication protocols based on group communication primitives. In *the International Conference on Distributed Computing Systems*, pages 156–163, Amsterdam.
- Koren, G. and Shasha, D. (1995). Dover : An optimal on-line scheduling algorithm for overloaded uniprocessor real-time systems. In *SISAM J. Comput.*, volume 24(2), pages 318–339.
- Liu, C. and Leyland, J. (1973). Scheduling algorithms for multiprogrammig in hard real-time environment. In *Journal of the ACM*, volume 20(1).
- M. Xiong, K. Ramamritham, J. H. and Stankovic, J. A. (2002). Mirror: A state-conscious concurrency control protocol for replicated real-time databases. In *Information systems*, volume 27(4), pages 277–297. Elsevier Science Publishers.
- Mok, A. M. and Chen, D. (1997). A multiframe model for real-time tasks. In *IEEE transactions on Software Engineering*, volume 23(10), pages 635–645.
- Pedone, F., Guerraoui, R., and Schiper, A. (1998). Exploiting atomic broadcast in replicated databases. In *the 4th International Euro-PAR'98 Conference*, pages 514–520, Southampton. Lecture Notes in Computer Science 1470.
- Son, S. (1987). Using replication for high performance database support in distributed real-time systems. In *the 8th IEEE Real-Time Systems Symposium*, pages 79–86.
- Son, S. and Kouloumbis, S. (1993). A real-time synchronization scheme for replicated data in distributed database system. In *Information Systems*, volume 18(6).
- Son, S. and Zhang, F. (1995). Real-time replication control for distributed database systems: Algorithms and their performances. In *the 4th International Conference on Database Systems for Advanced Applications*, Singapore.
- Ulusoy, O. (1994). Processing real-time transactions in a replicated database system. In *Distributed and Parallel Databases*, volume 2, pages 405–436.

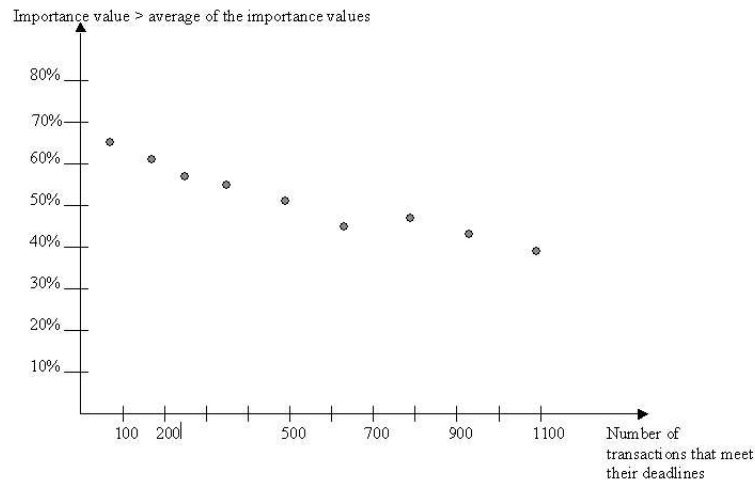


Figure 4: The percentage of transactions (with high importance values) that meet their deadlines when overloading the system