Management of a Remote Backup Copy for Disaster Recovery

RICHARD P. KING and NAGUI HALIM
IBM T. J. Watson Research Center

and

HECTOR GARCIA-MOLINA and CHRISTOS A. POLYZOIS
Princeton University

A remote backup database system tracks the state of a primary system, taking over transaction processing when disaster hits the primary site. The primary and backup sites are physically isolated so that failures at one site are unlikely to propagate to the other For correctness, the execution schedule at the backup must be equivalent to that at the primary When the primary and backup sites contain a single processor, it is easy to achieve this property However, this is harder to do when each site contains multiple processors and sites are connected via multiple communication lines We present an efficient transaction processing mechanism for multiprocessor systems that guarantees this and other important properties. We also present a database initialization algorithm that copies the database to a backup site while transactions are being processed.

Categories and Subject Descriptors: C 2.4 [**Computer-Communication Networks**]: Distributed Systems—*distributed applications, distributed databases*; D.4 5 [**Operating Systems**]: Reliability—*backup procedures, fault tolerance*; H.2.2 [**Database Management**]: Physical Design—*recovery and restart*; H.2.4 [**Database Management**]: Systems—*concurrency, distributed systems, transaction processing*; H 2.7 [**Database Management**]: Database Administration—*logging and recovery*

General Terms: Algorithms, Reliability

Additional Key Words and Phrases: Database initialization, hot spare, hot standby, remote backup

1. INTRODUCTION

In critical database applications, the halting of the computer system in case of failure is considered unacceptable. Instead, it is desirable to keep an up-to-date backup copy at a remote site, so that the backup site can take over

Authors' addresses: R. P. King and N. Halim, IBM T. J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598; H. Garcia-Molina and C A. Polyzois, Dept of Computer Science, Princeton Univ , Princeton, NJ 08544.

transaction processing until the primary site recovers. Such a remote backup database (or *hot standby* or *hot spare*), normally of the same capacity as the primary database, should be able to take over processing immediately. Furthermore, when the primary site recovers, the backup should provide it with a valid version of the database that will reflect the changes made by transactions processed while the primary was not operational; this will enable the primary site to resume normal processing. Finally, the overhead at the primary site for the maintenance of the remote backup copy should be kept low.

A remote backup copy has advantages over other forms of fault tolerance. To illustrate, consider a common form of local replication: mirrored disks. It can happen that a repairman trying to fix one of a pair of mirrored disks accidentally damages the good disk, which is physically located next to its faulty mirror image [8]. Thus, because failures tend to propagate, local replication may sometimes be inadequate. The remote backup-copy technique decouples the systems physically, so that failures are isolated and the overall system is more reliable. We have mentioned only hardware failures so far; physical isolation may also contain *some* failures caused by operator errors or software bugs. For example, an operator may mistakenly destroy the database by reformatting the disks that hold it. This has actually been reported [8]. However, it will be much harder for the operator to destroy the database stored remotely and under the control of a separate operator. Similarly, software bugs triggered by particular timing events at the primary site will probably not occur at the backup. The backup will have bugs of its own, but these are likely to occur at different times. Finally, environmental disasters (e.g., earthquakes, fires, floods, power outages) are yet another important category of disasters whose impact can be minimized through the use of a remote backup. Thus, remote backup copies provide a relatively high degree of failure isolation and data availability, and are actually used in practice [7].

The backup may find other uses as well, e.g., database schema or software updates can be performed without interrupting transaction processing, by having the backup take over processing while the new schema/software is being installed at the primary and then reversing the roles to update the backup.

Backup systems can track the primary copy with varying degrees of consistency. An *order-preserving* backup ensures that transactions are executed in the same logical order they were run at the primary. This is the only approach we consider here. Nonorder-preserving backups are sometimes used in practice [5], but may lead to inconsistencies between the primary and the backup.

Along another dimension, backup systems can run *1-safe* or *2-safe* transactions [10, 13]. 2-safe transactions are atomic: either their updates are reflected at both the primary and the backup, or they are not executed at all. A conventional two-phase commit protocol can be used to provide 2-safety [9, 16]. One major drawback of 2-safe protocols is that they increase transaction response time by at least one primary-backup round trip delay plus some processing time at the backup. According to Lyon [13], this may exceed one

second in practice. These delays force transactions to hold resources (e.g., locks, workspace, etc.) longer, thus increasing contention and decreasing throughput.

To avoid these delays, many applications use 1-safety only: transactions first commit at the primary and then are propagated to the backup. A disaster can cause some committed transactions to be lost. These losses only occur when a disaster hits and are "economically acceptable" in applications with "very high volumes of transactions with stringent response time requirements. Typical applications include ATM networks and airline reservations systems [13]." The backup algorithm we present in this paper is intended for very high performance applications and only provides 1-safety.

Apart from reducing contention for resources, 1-safe transactions have some other advantages as well. Suppose we have a system with some response-time requirement. Typically, this may be something like "90% of the transactions must have response time below t_r." Assume that given this requirement and using 2-safety the system can achieve a maximum throughput w. If we switch from 2-safety to 1-safety, the response time will drop for all transactions. Consequently, we can increase the load of the system beyond w and still meet the response time requirement. Thus, we can trade the gain in response time for an increase in throughput.

A third advantage of 1-safety is the simplification of processing when the backup becomes unreachable. With 2-safety, the primary site must change the way it processes transactions to "skip" the agreement phase. When the backup becomes reachable again, it must catch up with the primary (in a special processing mode). Then the two sites must synchronize, revert to the previous mode and resume normal processing. 1-safety makes things much easier: if the backup becomes unreachable, the messages are simply accumulated at the primary and are later sent to the backup. No deviation from normal processing occurs. This is especially convenient for short periods of disrupted communication (longer failures of communication links may require reinitialization of the backup anyway, since the backlog of messages may have grown too big).

Finally, with 1-safety it is easier to support multiple backups of the same database than it is with 2-safety. When 2-safety is used, the coordinator of a transaction must wait for messages to be received from all (or at least the majority) of the participants. Thus, the latest response "sets the pace." The more sites there are, the more likely it becomes that one of them will introduce some delay and cause the other sites to wait for it. Furthermore, when the configuration changes (sites leave or join the set of active sites), all sites have to be notified. The problems mentioned above do not appear under 1-safety: each site operates at its own pace and independently of any other site.

We would like to emphasize that we are not excluding the possibility of 2-safe transactions. As a matter of fact, within the same application it is possible and useful to run some transactions as 1-safe and others as 2-safe. For example, in a banking application, transactions involving large amounts of money could be run as 2-safe, while the bulk of the transactions, involving

relatively small amounts, can be run as 1-safe. The 1-safe backup algorithm we present can be combined with a 2-safe one in this way, although we do not show this here.

In this paper we make three main contributions:

(1) We precisely define the concept of 1-safe transactions and its implications. Commercial systems claim they provide 1-safe transactions, but they never state precisely what it means to "lose some transactions." For instance, it is important to specify that if a transaction is lost, any transactions that depend on it cannot be processed at the backup.
(2) We present a fully decentralized and scalable backup management algorithm that does not rely on a single control process. All commercial products concentrate the logs of every transaction into a single control process and master log. Such a process may become a bottleneck in larger systems. Furthermore, their solutions are not amenable to parallelization: if the control process is split into a set of processes, correctness can no longer be guaranteed.
(3) We present an efficient database initialization algorithm. It does not rely on first making a fuzzy dump and then bringing it up-to-date with the log. Instead, the fuzzy dump and the log playback occur concurrently, making the initialization simpler and faster.

The paper consists of two parts. In Sections 2–4 we give the necessary background (previous work, architectural model, goals) and in Sections 5–10 we present our solution.

2. REVIEW OF EXISTING SYSTEMS

Systems for maintaining a backup copy are commercially available. For example, Tandem provides a Remote Duplicate Database Facility (RDF) [18] and IBM markets an Extended Recovery Facility (XRF) [12]. The latter is primarily appropriate for local backups. There is also a research project at IBM [3], which is intended to support remote backups. These packages provide a set of utilities for dumping databases, monitoring them, and propagating modifications to a backup database. It is not our intention here to describe the full packages; we are only interested in the algorithms used for maintaining and initializing the backup database.

We discuss RDF briefly, which is typical of commercial systems. At the primary site, undo/redo log entries for every transaction are written on a master log. (If several processes or processors are executing transactions, they all communicate with a logging process that manages the master log.) As this log is written, a copy is sent to a *control process* at the backup site.

When a transaction commits at the primary site, it is assigned a *ticket* (a sequence number) that represents its commit order relative to other transactions. This is the order in which transactions must install their updates at the backup. Note that the notion of a ticket may not be explicitly implemented; it may be implicit in the position of a commit message in the log. The log entries are received and installed at the backup by the control

process in the same order they were generated. When a commit record for a transaction T_a is encountered, all transactions with earlier tickets have already been safely received at the backup, so it is safe to commit T_a.

The actual writes are performed by a collection of *writing processes.* Each writing process is assigned a part of the database (e.g., a disk volume). The control process distributes the writes to the appropriate writing processes, which install the updates in the order they receive them.

To initialize the backup database, a fuzzy dump is made at the primary, usually onto tape. While the dump is in progress, the master log is sent to the backup, where it is accumulated on disk. The tape is carried over to the backup (or transmitted if it is small) and loaded onto the database. Finally, the saved log is played back against the database, bringing it up to date. The log playback is done through the control process and ticketing described above.

The ticket assignment process plays a critical role, since every transaction has to go through it. Having this central bottleneck is undesirable in multiprocessor systems. It would be much better to have multiple master logs created and received by multiple control processes running on multiple machines at each site. Unfortunately, it is now much harder to know when a transaction can be committed at the backup and what its ticket is. But before we can illustrate this difficulty and our solution, we must step back and define more formally an architectural framework and correctness criteria.

3 ARCHITECTURE AND GOALS

To make our discussion concrete we define an architecture, shown in Figure 1. We try to place as few restrictions as possible on this model, in order to make it widely applicable. By *site* we mean all of the computer equipment at one of the locations where the database resides. Each site (primary and backup) consists of a number of *stores* and a number of *hosts*. The data actually resides in the stores, which are assumed to have reliable, nonvolatile storage. The stores also have enough processing power to perform basic database operations (such as accessing and modifying records), to keep logs, etc. The hosts are responsible for processing the transactions; they communicate with the users, the stores at the local site and the hosts at the remote site. We assume a fail-stop model [17] for host and store processors.

The stores and hosts need not necessarily be disjoint processors. The same processor could perform both tasks, by dividing its CPU cycles between them. Stores and hosts might actually be implemented as two layers of software in the same system. For example, Bernstein et al. [2] define a database system architecture with a *transaction manager* (corresponding to our notion of host) and a *data manager* (corresponding to our notion of store). Similarly, the RSS store manager of system R could implement our stores and the RDS system our hosts. The problems and algorithms presented in the following sections generally apply to both cases, i.e., when stores and hosts are disjoint processors and when a computer acts partly as a store and partly as a host. However, when we consider single failures in Section 9, we have to distinguish between the two models.

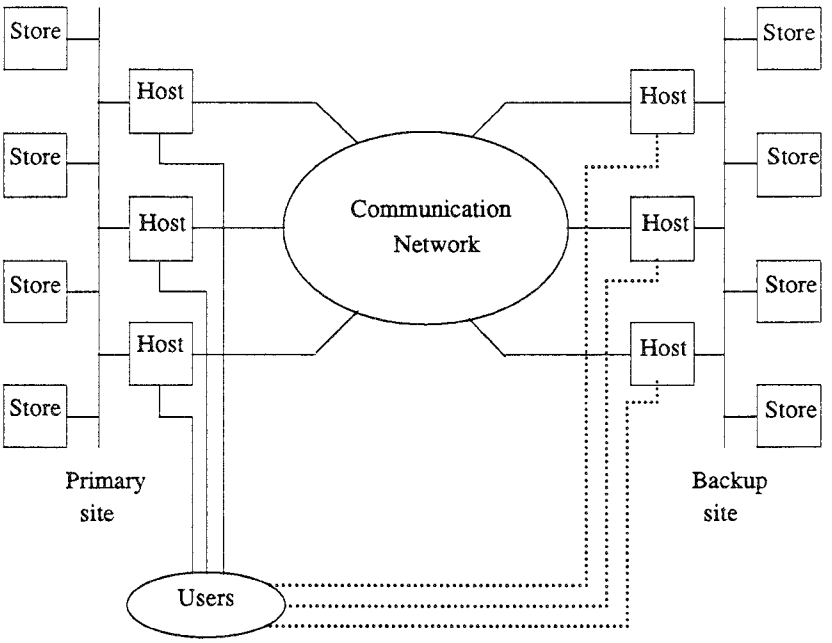

Fig. 1. System architecture.

All of the stores and hosts at one site can communicate with each other through shared memory or through a local network or bus. The method we present applies to shared memory architectures (e.g., DEC Firefly, Encore Multi-Max) as well as more loosely coupled systems (e.g., a VaxCluster, a Tandem Dynabus or a set of Camelot workstations connected via an Ethernet).

Running between the two sites are several communication lines, which let the hosts at the primary site send copies of operations being performed to the backup hosts. Control messages are also exchanged over these lines. The connections between primary and backup hosts may be either of the datagram or of the virtual circuit type [19]. In the following sections the type of the connections usually does not make a difference. When it does (e.g., when preservation of the order of messages is significant), we mention so explicitly. The existence of enough bandwidth to propagate the changes is assumed; however, communication delay is not a critical issue here, since we are assuming 1-safe transactions.

The database model we use is a simple version of the relational model, but it is realistic enough to allow us to study the major issues. The database contains a set of *tables*, and each table contains a set of *records*. The tables have unique *names* and the records have unique *record ids*. Requests can be made to create or delete tables and to insert, select, update or delete records. Each of these requests must provide the appropriate parameters. For example, in order to update a record, one must provide the id of that record along

with the new version of it. The store will, upon request, create and maintain indices on tables, with any subset of the fields forming the key of the index. Basic transaction processing functions, such as locking, abort and commit operations, are also supported by the store. Such requests are always associated with a transaction identifier, which is established with a begin-transaction request during the initialization phase of a transaction.

The tables are partitioned among the stores. This could be done by hashing names, by table look-up or by some other method; the details are not important, but as we show in Section 10, the partition must be *logically* identical at the two sites. We also remind the reader that the capacity of the stores at the remote site must be at least equal to that of the stores at the primary site.

Let us take a brief look at how normal processing would proceed at the primary site without the existence of the backup. A host gets a transaction from a user and assigns it a transaction id (host id followed by sequence number). Before a transaction issues the first request to a particular store, it must send a begin-transaction request to this store. Then the transaction is executed, by issuing the appropriate requests to the store(s) where the corresponding data reside; the requests contain the transaction id and the necessary parameters. When all of the requests have been executed, the host initiates a two-phase commit protocol that ensures that the transaction is either committed or aborted at all stores. The stores, on the other hand, execute all of the database operations that are issued to them, produce the necessary logs, set the appropriate locks, etc., [9]. We are assuming strict two-phase locking is used for concurrency control. (This is not necessary. In Section 10 we drop this assumption and show that our mechanism also applies to systems that do not use two-phase locking.) Note also that no global transaction sequence numbers indicating the commit order are generated. (Generating them would create a bottleneck.)

As failures occur at the primary, the system tries to recover and reconfigure. (The details for this are given in Section 9.) However, multiple failures may slow down the primary site or even stop it entirely. At this point, a *primary disaster* is declared and the backup attempts to take over transaction processing. The declaration of a disaster will in all likelihood be done by a human administrator. This is mainly because it is very hard for the backup site to distinguish between a catastrophic failure at the primary and a break in the communication lines. User terminals are usually connected to both the primary and the backup site. The backup connection is normally on standby and is used in case of disaster to route input transactions to the backup. For simplicity we assume that when a primary disaster occurs, the hardware at the backup site is fully operational. (Certain failures at the backup could be tolerated during a primary disaster, but this is not discussed here.) A *backup disaster* is similarly declared when failures impair the backup site. We assume that the primary site is operational during a backup disaster.

A site is always in one of three modes: primary, backup or recovering. At most one of the sites can be in the primary mode at any time. Under normal operation, processing of the transactions takes place at the site operating in

primary mode and sufficient information is sent to the site operating in backup mode to allow it to install the changes made at the primary site. When a primary disaster is declared, the site previously operating in backup mode starts operating in primary mode and all of the transactions are directed to it. When the failed site comes up, it enters a special recovering mode, which will allow it to get a consistent, up-to-date copy of the database and (perhaps later) resume normal processing. The recovering mode is also used for creating the backup copy at system initialization. Note that the primary and backup roles are interchangeable between the two sites.

Our remote backup algorithm was designed with the following goals in mind:

Database consistency. We require that the database at the backup be up-to-date and consistent. Application programs are usually written under the assumption that they will operate on correct data, and error handling in these programs is rarely comprehensive. In case of disaster the backup takes over transaction processing and, if the backup does not preserve consistency, the application programs may have to run on an inconsistent copy. Such an inconsistency could lead to delays in transaction processing or even to crashes. Thus, compromising the consistency of the database may endanger its continuous operation and should be avoided. In Section 4 we define the consistency requirements in more detail.

Scalability. Most existing systems have a component which must "see" (process in some way) all transactions. For example, the master log may be such a component. As systems scale upwards and use multiple computers for transaction processing, the performance of this component will eventually form a bottleneck in the overall performance of the system, no matter how small the amount of processing for each transaction is. To illustrate, consider the master log and assume that all messages sent to the backup must go through the same "concentrating" processor that is connected to the communication line. Suppose that this processor has a capacity of P instructions/sec and that processing a message (e.g., a message containing log or control information from a primary to a backup computer) requires x instructions. If a transaction generates m messages on the average and there are N processors processing transactions, the logging process limits the throughput of each processor to $P/(x \times m \times N)$ transactions per second, no matter what the bandwidth of the communication link is.

Parallelism at the backup. Although it is rather difficult to give a measure for parallelism, one can see that an acceptable solution to the problem should exploit the potential for parallelism at the backup. For example, suppose that transactions at the backup site are processed sequentially. If the primary site employs multiprogramming and/or multiprocessing to allow parallel execution of transactions, it will have a higher throughput than the backup. The backup will inevitably be unable to keep pace with the primary, and the backup copy will become out of date.

Primary overhead minimization. We try to minimize the overhead induced by the backup algorithm at the primary site. During normal processing, the backup is simply installing updates, as opposed to the primary, which is actually running the transactions. Yet, the backup should be capable of processing transactions after a disaster, so we expect it to have some spare capacity during normal processing. Thus, if we have an option of doing some backup function either at the primary or at the backup, we opt for doing it at the backup. This lets the primary dedicate its resources to transaction processing.

We note that our goals do not necessarily hold in every case. For example, the backup may not have spare capacity during normal processing (e.g., if it is being used for some other, noncritical processing), or database consistency may not be significant. However, we do believe that the stated goals are desirable in an important subset of applications, so that it is worth investigating a solution for this scenario.

4. CORRECTNESS CRITERIA

In this section we describe the database consistency requirements for 1-safe transactions more precisely. The transaction processing mechanism at the primary site ensures that the execution schedule PS of a set of transactions T is equivalent to some serial schedule. Schedule PS induces a set of dependencies on the transactions in T. Let T_x and T_y be two transactions such that T_x commits before T_y. We say $T_x \rightarrow T_y$ (in PS) if both transactions access a common data item and at least one of them writes it [11]. Dependencies can be classified into write-write (W-W), write-read (W-R) and read-write (R-W) depending on the actions that generate them.

The backup site will execute a subset of the actions in PS. Let this schedule be BS. Read actions are not performed at the backup since they do not alter the database. The write actions that are executed simply install in the database the value that their counterpart installed at the primary. Because of failures, not all write actions of PS may appear in BS.

Requirement 1: *Atomicity.* If one action of a transaction T_x appears in BS, then all write actions of T_x appearing in PS must appear in BS. This disallows partial propagation of a transaction. Let R be the set of transactions whose writes are in BS, $R \subseteq T$.

Requirement 2: *Mutual consistency.* Assume T_x and T_y are in R. Then, if $T_x \rightarrow T_y$ in BS, it must be the case that $T_x \rightarrow T_y$ in PS. This guarantees that the backup schedule is "equivalent" to the primary, at least as far as the propagated *write* actions are involved. (Since no read actions take place at the backup, this requirement does not apply to R-W or W-R dependencies.)

Let M be the set of transactions that were not fully propagated to the backup before a failure and hence were not installed. In addition to these transactions, there may be other transactions that we do not want to install at the backup. For example, suppose that when T_x and T_y execute at the primary, T_x writes a value read by T_y. If T_x is not received at the backup

(i.e., $T_x \in M$), we do not want to install T_y either, even if it is properly received. If we did, the database would be inconsistent.

To illustrate this, say that T_x is the transaction that sells a ticket to an airline customer. It inserts a record giving the customer's name, date, flight involved, payment information, and so on. Transaction T_y checks-in the passenger at the airport, issuing a seat assignment. The updates produced by T_y cannot be installed at the backup without those of T_x: there would be no passenger record to update. Thus, we have the following requirement:

Requirement 3: *Local consistency.* No transaction in R should depend on a transaction in M. That is, suppose there is a transaction $T_a \in M$ and there is a sequence of W-W and W-R dependencies (*not* R-W) in PS:

$$T_a \rightarrow T_b \rightarrow T_c \rightarrow \cdots \rightarrow T_n.$$

Then none of $T_a, T_b, \ldots, T_n$ is allowed to be in R. If C is the set of transactions that depend in this way on M transactions, then $R \cap (M \cup C)$ should be empty.

At this point we would like to make two observations. First, R-W dependencies do not cause violations of the local consistency constraint. If $T_a \rightarrow T_b$ and the only dependencies between these two transactions are R-W, then the values installed by T_a cannot possibly affect the values installed by T_b. Thus, one can install at the backup the updates made by T_b and have a consistent database, even if T_a does not reach the backup. Our second observation is that the notion of local consistency is similar to that of recoverability as defined by Bernstein et al. [2]. Since we are also allowing W-W dependencies (not just W-R) in the dependency path of the definition above, Bernstein et al. define local consistency comes closer to the concept of strict executions. The motivation, however, is different in the two cases. Bernstein et al. [2] introduced strict executions to avoid problems associated with uncommitted data being read or overwritten by another transaction and are intended to achieve correctness of a schedule within a *single* site. In our case, local consistency is required to ensure the equivalence of *two* schedules. It only happens that both problems are dealt with in the same way: the actions of some transactions are delayed until some other transactions have committed. Since the two notions are quite different in context, we have chosen a distinct name for our requirement.

Requirement 4: *Minimum divergence.* The backup copy should be as close to the primary as possible, i.e., the backup site should commit as many as possible of the transactions it is given, as long as none of the above correctness constraints is violated. In other words, if a received (therefore not belonging to M) transaction does not depend on any transaction in M (i.e., does not belong to C), then it has to belong to R, i.e., $T = R \cup M \cup C$.

In closing this section we make four observations. First, we have implicitly assumed that the primary and backup schedules PS and BS run on the same initial database state. In Section 8 we present a procedure for initializing the backup database so that this property holds.

Our second observation is that read-only transactions do not modify the state of the database and therefore need not be propagated to the backup site.

The third observation is that most replicated data mechanisms described in the literature (e.g., Bernstein et al. [2], Garcia-Molina and Abbott [6]) would not allow what we call missing transactions. That is, they would have the property $M = \emptyset$ (2-safe transactions). As discussed in the introduction, they would use a two-phase commit protocol to achieve this property. With the weaker constraints we have defined in this section, it is possible to have a more efficient algorithm. In essence, transactions commit at the primary (using a local two-phase commit that involves primary hosts and stores), release their locks, and only then are propagated to the remote backup. At the backup there will be a second local commit protocol to ensure that actions are properly installed there. The details are given in the rest of the paper.

Our last observation deals with the semantics of missing transactions. In particular, is it "valid" to lose transactions that already committed at the primary? Is it "reasonable" to throw away transactions that were propagated but violate the local consistency requirement? As discussed in the introduction, real applications do allow missing transactions and can cope with the consequences. A simple example may illustrate what these consequences are.

Consider a banking application where a transaction T_x deposits 1,000 dollars into an account that initially has no funds, and a transaction T_y withdraws 400 dollars from the same account. Both of these transactions run at the primary just before a disaster. Transaction T_x is lost, but T_y arrives at the backup. (This can happen if the log records for T_x are propagated via a different communication line than that for T_y.) To satisfy the local consistency constraint, T_y is not installed at the backup. The database will be consistent at the backup but not consistent with the real world. The inconsistencies with the real world can be detected and corrected manually. For example, when the customer checks the balance of the account (at the backup) and sees zero instead of 600, he will go to the bank with his deposit receipt and ask for a correction. The bank, knowing that a disaster occurred, will be willing to make amends for missing transactions. The bank might lose some money (especially if withdrawal transactions are lost), but this cost can be much smaller than the cost of providing 2-safe transactions. Furthermore, not all transactions have to be 1-safe. As noted in the introduction, transactions that involve large sums of money can be performed with 2-safe protocols.

Finally, there is the issue of whether T_y (the withdrawal transaction that depends on the lost T_x) should be discarded even if it arrived at the backup. Installing T_y makes the database inconsistent (e.g., there are not enough funds in the account to cover this withdrawal). On the other hand, discarding T_y creates an inconsistency with the real world. We believe it is more important to maintain database consistency. As discussed in Section 3, without consistency it would be very hard to continue processing new transactions. Furthermore, T_y can be saved for special manual processing. A bank employee can look at the T_y record and decide on the best compensating transaction to run. Since disasters occur rarely, the price to pay for this

special processing of the relatively few transactions that are discarded is probably acceptable.

5. GENERATING THE LOGS

In the remaining sections we present our solution. Our method is based on logs, which are generated at the primary site and propagated to the backup (this is discussed in the rest of this section). At the backup, two-phase locking is used to detect and reconstruct dependencies, in order to install the changes in a way observing the consistency requirements of Section 4. The details for this are given in Section 7. In Section 8 we present a method for initializing the database and in Section 9 we discuss partial failures. Section 10 contains some extensions to the basic algorithm and Section 11 concludes the paper.

The basic idea in our solution is to reproduce the actions of each primary store at a corresponding backup store. For this, a log of the executed actions must be forwarded to the backup store. The log could be at the action or the transaction level. We have chosen the former. The log could have undo/redo information, or just redo. Our method does *not* sent undo information; this reduces communication traffic. We made this set of choices for concreteness and for compatibility with our architecture; other alternatives may also be possible.

Given these design choices, we now describe in more detail how the primary site generates the logs. The stores at the primary site keep track in a redo log of the actions they are performing on behalf of a transaction. (As is explained later, read sets also have to be recorded.) Each redo log entry should contain enough information to enable the corresponding stores at the backup site to repeat the same operations on their copy. This information should contain the following data:

S_i:	the store performing the operation
H_j:	the host controlling the transaction
T_x:	the transaction id
act:	action descriptor (e.g., update)
tbl:	the name of the table involved
key:	the record id of the record being updated
val:	the after image of the record being updated

Note that not all entries contain all of the above data. For example, the log entry for a delete need not contain an after image, and for create-table no key nor value is necessary.

When a host receives a transaction, it issues a begin-transaction request to the stores (at the primary site) that will participate in the processing of the transaction. Processing proceeds as usual, with local locks being used for concurrency control and with the stores constructing the redo log. Note that

in many cases stores keep a log anyway (for local crash recovery), so the log for the remote backup represents no extra overhead. An alternative not considered here is for hosts to directly construct the log as actions are acknowledged from the stores.

When a transaction completes, the host initiates a (local) two-phase commit protocol. Upon receipt of a *commit* message, a store produces a local ticket number by atomically incrementing a counter and getting its new value. Intuitively, if a transaction gets a ticket t at store S_i, the transaction "saw" state $t - 1$ of the store. If the transaction has only read (not written) data at that store, the ticket of the transaction becomes the value of the counter plus one, but the counter is *not* incremented (since the state of the store did not change). For example, if the counter has the value 25, the committing transaction gets ticket 26. If the transaction wrote at this store, the counter becomes 26; otherwise, it stays at 25.

The store creates a commit entry in its redo log, with the following data:

S_i:	identifies the store involved
H_j:	identifies the host controlling this transaction
T_x:	transaction id
act:	action descriptor, in this case "commit"
ticket:	the local ticket number

Then the stores send an acknowledgement to the coordinating host and release the locks they held on behalf of the transaction.

Now the logs must be propagated to the backup. Stores typically do not run remote communication software, so the logs must be sent via the hosts. This can be done in two ways.

Bundles. The first way to send logs, called the bundle model, is on a per transaction basis. The coordinating host collects the log entries for the transaction (including the commit entries and the tickets) from the participating stores, bundles them together and sends them across the network to a remote host. The choice of the remote host may be static (i.e., host i at the primary site always sends to host j at the backup site) or it may be determined dynamically, perhaps on the basis of load information. The remote host receives the transaction's logs, unbundles them, partitions them by store and sends them to the appropriate (local) stores for processing.

Log streams. The second way to propagate logs to the backup, called the stream model, is on a per store basis. For each store, a local host is designated as its supervising host. A store communicates with its remote peer via the corresponding supervising hosts, i.e., the primary store sends a message to its supervising host, the supervising host sends it across the network to the supervising host of the backup store, which in turn delivers the message to the backup store. Thus, a log stream is formed from every primary store to its peer at the backup to propagate the log entries created by

the primary store. Note that under this model, parts of a transaction that executed at different stores may follow different paths to the backup site.

In our discussion we distinguish between the two models for log propagation when necessary. For example, as we see in Section 7, the satisfaction of the minimum divergence requirement depends on the model we adopt.

Each transaction id may be associated with several ticket numbers, one for each store in which actions of this transaction were executed. As is explained in the next section, these ticket numbers are used to ensure that the database changes of different transactions are applied to the backup copy in the same order as in the primary. After the processing of the transaction at the remote site has finished, the host that controlled the transaction at the primary site is informed, which in turn informs the stores that they may safely erase the relevant portion of the redo log from their storage. Note that, from the viewpoint of the primary site, transaction processing is similar to what it would be without a backup copy. The only difference is the *local* generation of a ticket number and the forwarding of log records to a remote host. We believe these activities do not impose a significant overhead on normal transaction processing, which is consistent with our goals. The reader should contrast our model with other models for parallel logs, e.g., Agrawal [1], where multiple logs are allowed, but all transactions have to go through a central processor referred to as the *back-end controller*.

There are two alternatives with respect to when the user submitting a transaction gets a response. If the response is sent to the user after the transaction has committed at the primary site, then, in case of disaster, the transaction may be lost if it does not commit at the backup. If the user gets a response after the transaction has committed at the backup, then it is guaranteed that the effects of the transaction will not be lost in case of disaster. Note that transactions are 1-safe in both cases; only the user's information about the fate of a transaction is different.

6. INSTALLING ACTIONS AT THE BACKUP—WHY THE SIMPLE APPROACH DOES NOT WORK

The next problem we address is installing the actions at the backup site. At this point, it is tempting to propose a very simple solution: as each backup store receives the redo log entries, it simply installs them in ticket order, without regard to what other backup stores may be doing. This simple solution could be viewed as the generalization of the approach used by commercial systems for the case where there is no master log.

In this section we show that there are problems with this approach. But before doing this, we point out that at least the simple approach does guarantee mutual consistency (Requirement 2). To see this, suppose we have two transactions T_x and T_y, such that $T_x \rightarrow T_y$ at the backup site. Let z be a data item that caused this dependency. If actions are installed at the backup in local ticket order, the ticket number of T_x is smaller than that of T_y at the particular store. This implies that at the primary T_x got its ticket before T_y. When T_x got its ticket, it held a lock on z, which was not released until T_x

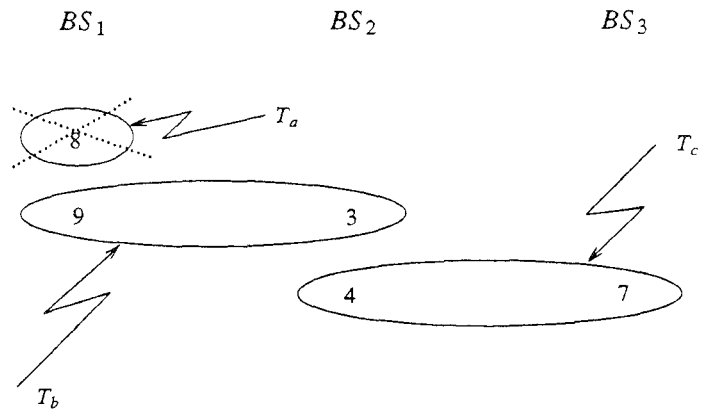

Fig. 2 Cascading aborts (Bundle model).

committed. The lock was incompatible with the lock on z requested by T_y. Thus, $T_x \rightarrow T_y$ at the primary.

One problem with the simple approach is that it does not ensure that transactions are installed atomically at the backup (Requirement 1). Thus, in addition to executing writes in ticket order at each store, it is necessary to execute a *local* two-phase commit protocol among the stores that participated in a transaction.

Even with installation in ticket order and a two-phase commit protocol, local consistency (Requirement 3) may be violated. To illustrate this, we need an example with three transactions and three stores (Figure 2). We assume the logs are propagated to the backup on a per transaction basis (as bundles). At the primary site, the following events occur: T_a writes data at store S_1 getting ticket number 8 (S_1). (We follow the ticket number by the store id to prevent confusion.) Transaction T_b then reads this data, getting ticket 9 (S_1) and writes at store S_2, receiving ticket 3 (S_2). Later on, T_c reads this data and writes at both S_2 (ticket 4) and S_3 (ticket 7).

Assume that each transaction is coordinated by a *different* host, so that the three transactions each follow a different path to the backup. A disaster hits and T_a does not reach the backup site. Transactions T_b and T_c do make it to the backup (Figure 2). At backup stores BS_2 and BS_3 all writes are received, so there are no gaps in the sequence numbers. Thus, it would appear that a transaction like T_c could commit: all of the data it wrote is available at BS_2 and BS_3 (i.e., the two-phase commit could succeed). Furthermore, all data written at those stores before T_c is also available (i.e., installation in ticket order is feasible). Unfortunately, since T_a is lost, at BS_1 there is a gap: sequence number 8 (S_1) is missing. Thus, T_b with sequence number 9 (S_1) must be aborted at store BS_1 and consequently (by the atomicity requirement) at all backup stores. If this happens, then there will be a gap at BS_2: sequence number 3 (for T_b) is not really available. Hence, writes with higher sequence numbers cannot be installed, so T_c must also be aborted.

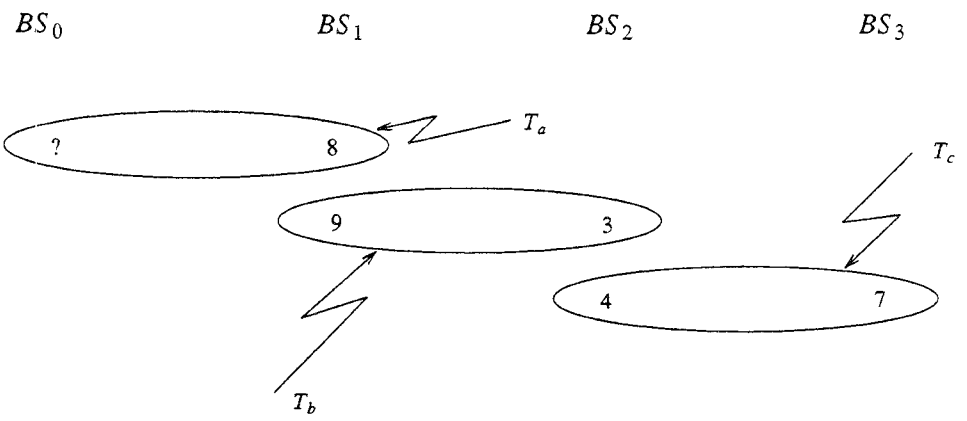

Fig. 3. Cascading aborts (Stream model).

The difficulty caused by this type of cascading aborts should be apparent by now: before a transaction can be installed at the backup, we must make sure that the transactions it depends on have *committed* at the backup. This involves synchronization with other stores. Thus, updates cannot be installed blindly by each backup store, even if they are in correct sequence.

An interesting point about cascading aborts is that they appear under both models for log transmission mentioned in Section 5. To illustrate the problem under the log stream model, consider the following example (Figure 3). Transaction T_a executed at stores S_0 and S_1. The logs created at S_1 reach the backup store BS_1, while the ones created at S_0 do not reach BS_0 (they were among the last few messages transmitted on the log stream for S_0 which were lost because of a disaster). The atomicity requirement prevents transaction T_a from committing. The rest of the example is the same as above and the effect of cascading aborts equally annoying. The persistence of this effect under two quite diverse models has led us to believe that this problem is inherent in the presence of multiple independent logs, i.e., when no *total* ordering of messages is present. We would like to stress the fact that the scenarios mentioned above are still possible even if the order of messages is preserved within each log stream (e.g., virtual circuit type connections).

As we have seen, installing actions in ticket order at each store guarantees mutual consistency. Unfortunately, there is potentially a major drawback in terms of efficiency. Let us illustrate with an example. Suppose that T_1 and T_2 access disjoint data sets at some primary store and $ticket(T_1) = ticket(T_2) - 1$. At the backup, the writes for T_2 cannot be installed until those for T_1 are installed (this is what ticket order means). Thus, T_2 must wait until T_1 *commits* (which, as we have seen, involves waiting for other stores to tell us what they have done), and then wait further until the writes of T_1 are actually executed. This is inefficient, especially if stores have a capacity for executing writes in parallel, e.g., have multiple disks. Even with one disk, efficient disk scheduling of T_1's and T_2's writes is not possible.

Remember that T_1 and T_2 do not depend on each other, so their commits are independent and their writes could proceed in parallel.

To avoid these delays, it is necessary for each store to determine if transactions depend on each other or not. If T_1 and T_2 write a common item, then they are certainly interdependent and T_2 must wait for T_1 to commit and write. However, even if T_1 and T_2 write disjoint sets, there may still be a dependency! These dependencies can only be detected if read sets of transactions are also propagated to the backup. To illustrate, suppose that at a primary store T_1 wrote item y and that T_2 read y and wrote z. There is a dependency $T_1 \rightarrow T_2$, so T_2 cannot commit at the backup unless T_1 does (local consistency, Requirement 3). The corresponding backup store will not be able to detect the dependency from the write sets ($\{y\}$ and $\{z\}$). If the read sets are sent, the dependency will be seen, and T_2 can be delayed until T_1 finishes.

In summary, we have seen that actions cannot simply be installed in ticket order at the backup stores. The system must guarantee that a transaction only commits when all of the transactions that it depends on have committed at the backup. In addition, a transaction should not wait for transactions it does *not* depend on, something that would happen if actions were done in *strict* ticket order. The mechanism we describe in the following section achieves these goals.

Finally, note that sending undo/redo logs (as opposed to simple redo logs as we are assuming) does not really eliminate the problems we have sketched here. If undo logs are included, it is possible for backup stores to install updates in local ticket order disregarding the state of other stores. If later it turns out that some transaction does not commit, its updates can be undone. This may lead one to think that it is possible to avoid these problems (especially cascading aborts) and the processing overhead to handle them by deferring the commit decisions until disaster time and doing some extra processing at that time to abort those transactions that cannot commit. However, the commit decisions must still be made at some point, and still involve making sure that all transactions that a transaction depends on have committed. It is not a good idea to delay all commit decisions until a disaster hits, mainly for two reasons:

(1) The undo logs for all transactions must be kept until the commit decision is reached, since potentially any transaction can be affected by a cascading abort (determining that a transaction can no longer be affected by a cascading abort is equivalent to making a commit decision). The logs grow with time, and it will probably be impossible to keep them because of space limitations.

(2) Processing of new transactions at the backup after a disaster would be delayed until the commits complete. This may take a long time, since the cascading aborts mentioned above may lead to extensive searches in the logs.

Thus, even with undo logs, we would still need a commit protocol to run as transactions are received, and it would be similar to the one we describe for

redo logging only. As we have stated, redo logging sends less data to the backup, so in the rest of the paper we only deal with it.

7. INSTALLING ACTIONS AT THE BACKUP—OUR APPROACH

In this section we present our method for installing the redo logs at the backup, prove its correctness and discuss its performance. In what follows, the notations T_x and $T(s_x)$ are equivalent and are both used to denote the transaction with ticket number s_x. The store is usually implied by the context.

No real processing of transactions takes place at the backup site, in the sense that no computation is performed for a transaction. The backup site is given the actions that were performed on the primary copy (in the form of a redo log) and has only to install the changes in a manner consistent with our correctness criteria. A backup host does not have much work to do. In the bundle model, it gets a redo log from a primary host, partitions the log by store and sends each part to the appropriate store. This can be viewed as the first phase of a (local) two-phase commit protocol, for which the host acts as the coordinator. In the stream model, a backup host simply has to forward the log streams to the stores for which it has been designated supervising host. In the stream model a coordinating host must be elected for a transaction and agreed upon by all participating stores. In order to avoid expensive election protocols, the coordinating host at the backup can be selected by the coordinator at the primary when a transaction commits there and recorded in the commit log entries of the participating stores. Alternatively, the coordinator at the backup could be determined using a mapping mechanism from transaction id's to backup hosts.

The stores install the changes (in the proper order, see below), and when they are finished, they send an acknowledgement to the coordinating host. When that host receives acknowledgements from all of the stores, it executes the second phase of the commit protocol and also sends an acknowledgement to the primary host (as mentioned in Section 5). The two-phase commit protocol is used to make sure that the changes of a transaction are installed atomically, so that the backup copy is never left in an inconsistent state.

In Section 6 we saw that updates in strict sequence order reduce parallelism. The intuition behind our solution is to detect exactly those cases where waiting is necessary and to let all other cases take advantage of parallelism. This is achieved through locks on the data items accessed by the transactions, which are *granted to the transactions in ticket number order*. Write (exclusive) locks are requested for items that are to be updated. For other items in the read set, read (shared) locks are requested. Additionally, a read lock on every table "name" accessed is requested, in order to ensure that the table is not deleted while accessing one of its records; if the table was created or destroyed, this lock must be exclusive.

For a concrete example, suppose that T_x has a smaller ticket number than T_y at one of the stores. If they access a data item in conflicting modes, our mechanism ensures that the lock is granted to T_x, the transaction with the

smaller ticket number. Transaction T_y cannot get the lock until T_x releases it, i.e., until T_x commits. If, on the other hand, there is no dependency between the two transactions, then they will not ask for conflicting locks, so they will be able to proceed in parallel.

We now describe the locking mechanism at the backup in detail. When a redo log for transaction T_x with ticket number s_x arrives at backup store BS_j, it is placed in a queue of transactions ordered by ticket number. In this queue, T_x can be in a number of states:

LOCKING: the transaction arrives and requests locks for all of the records and tables it accesses. Then it waits until all transactions with smaller ticket numbers have entered (or gone past) the SUBSCRIBED state (so that conflicts can be detected). Only then does it enter the SUBSCRIBED state.

SUBSCRIBED: the transaction is waiting until it is granted the locks it requested. The transaction may then proceed to the PREPARED state.

PREPARED: an acknowledgement for the first phase of the two-phase commit protocol has been sent to the backup coordinating host.

COMMITTED: the message for the second phase has been received for this transaction. All updates have been made public and all locks have been released.

When T_x arrives, BS_j sets the state of T_x to LOCKING and starts requesting the locks required. Transaction T_x asks for a lock on data item z by inserting itself in a list of transactions asking for a lock on z; the list is sorted by ticket number. Each store has a counter that keeps track of the local ticket sequence, showing the ticket of the last transaction that entered the SUBSCRIBED state at this store.

The locking procedure is summarized in Figure 4. After all of the locks for T_x have been requested, T_x waits for the counter to reach $s_x - 1$, if it has not already done so. This recursively implies that T_x will wait for all transactions with smaller ticket numbers to enter the SUBSCRIBED state. (Note that some or all of these transactions may be further ahead, in the PREPARED or COMMITTED state. The important thing is that they must have reached at least the SUBSCRIBED state before T_x can do so.) When this happens, T_x enters the SUBSCRIBED state itself. If T_x writes data at this store, then it increments the counter by 1. The increment in the counter may in turn trigger $T(s_x + 1)$ to enter the SUBSCRIBED state and so on. For example, if the current value of the counter is 25, then if transaction T_x with ticket 26 is in the LOCKING state, it enters the SUBSCRIBED state and increments the counter to 26, which may cause the transaction with ticket 27 to enter the SUBSCRIBED state and so on. If T_x were read only for this store, it would proceed to the subscribed state without incrementing the counter to 26.

```
state(T_x) = locking;
t = ticket(T_x) at this store;
FOR each object z accessed by T_x
  add lock request (includes T_x's ticket) to z list;
WAIT UNTIL (counter > = t - 1);
IF T_x writes at this store THEN
  counter = counter + 1;
state(T_x) = subscribed;
FOR each object z accessed by T_x
  WAIT UNTIL T_x's request is at head of z list;
state(T_x) = prepared;
send acknowledgement to coordinating host;
WAIT UNTIL commit message for T_x is received;
FOR each object z accessed by T_x
  BEGIN
    IF T_x wrote z THEN
      install new z value;
    remove T_x lock request from z list;
  END;
```

Fig. 4. Store pseudo-code for backup transaction processing.

We would like to note that a transaction only waits for (parts of) other transactions that executed at the *same* store to enter the SUBSCRIBED state. The parts of a transaction at different stores proceed *independently* from each other (this also avoids deadlocks). The coordinating host ensures transaction atomicity by waiting for all parts of the transaction to enter the PREPARED state before issuing a commit message and allowing any part to enter the COMMITTED state. If one part is delayed in entering the PREPARED state, the other parts will have to wait. In case of disaster one or more parts may never reach the PREPARED state. In this case the transaction cannot commit, and the parts that did reach the PREPARED state will have to be rolled back.

Waiting for transactions with smaller ticket numbers to enter the SUBSCRIBED state is necessary in the bundle model, because transactions that executed at the same store may follow different paths to the backup and therefore arrive out of order. In the stream model the transactions arrive in ticket order at a backup store, but it may still be necessary for them to wait for the counter to reach the appropriate value. For example, if multiprocessing is used at the backup, the transactions may try to enter the SUBSCRIBED state out of ticket order because of delays introduced by scheduling or because one of them had to ask for more locks than the other.

In the SUBSCRIBED state T_x waits until all of its lock requests reach the head of their corresponding lists. (A read request is also considered at the head of a list if all requests with smaller ticket numbers are for read locks.) When this condition is met, T_x enters the PREPARED state and informs the coordinating host. After commit, all of T_x's requests are removed from the corresponding lists.

When a failure occurs at the primary site, the backup is informed that primary processing will be switched to it. The (former) backup tries to

commit all of the transactions that can commit and aborts the ones that cannot. It then enters the primary mode and takes over processing.

Observation: The atomicity constraint (Requirement 1) holds.

Argument: Atomicity is enforced by the local two-phase commit protocol at the backup.

Observation: The mutual consistency constraint (Requirement 2) holds.

Argument: Since locks are granted in ticket order at the stores, updates inducing dependencies are installed in ticket order. As discussed at the beginning of Section 6, mutual consistency is observed.

Observation: The local consistency constraint (Requirement 3) holds.

Argument: Suppose $T_a \rightarrow T_b \rightarrow T_c \rightarrow \cdots \rightarrow T_n$ and T_a has not arrived yet at the backup site (we remind the reader that the dependencies are non R-W). Further, assume that $T_a \rightarrow T_b$ occurs at store S_1, $T_b \rightarrow T_c$ at S_2, and so on (note that these stores are not necessarily different). Since the dependency is non R-W, the ticket number of T_b is *higher* than that of T_a. According to our processing rules, T_b cannot enter the SUBSCRIBED state at BS_1 and will be unable to commit. At BS_2, T_b will lock the objects that caused $T_b \rightarrow T_c$ and will not release them until it commits, thus preventing T_c from entering the PREPARED state. Transaction T_c in turn prevents T_d from committing and so on. Thus, local consistency holds.

Observation: The minimum divergence constraint (Requirement 4) holds in the stream model (with message order preservation).

Argument: Consider a transaction T_x whose parts have been properly received at the backup stores and which does not depend on any transactions that cannot commit at the backup. In the stream model (and assuming that the order of messages is preserved, e.g., virtual circuit type connections), all transactions with smaller ticket numbers at those stores have arrived, so that the corresponding transactions can ask for the locks. Since there is no gap in the ticket sequence, all of these transactions will enter the SUBSCRIBED state at these stores. Note that some of these transactions may (in case of disaster) never enter the committed state or even go beyond the SUBSCRIBED state (perhaps because of the cascading aborts effect or because they depend on transactions that cannot commit). The important point is that these transactions *can* enter the SUBSCRIBED state, which will allow T_x to also do so. Since T_x does not depend on any transaction that cannot commit, the locks it requests will be available or they will be eventually released (since the transactions that hold them will commit), so that T_x will be able to obtain its locks, install its changes and finally commit.

Unfortunately, under the bundle model, the mechanism does not always satisfy the minimum divergence constraint. Since transactions that executed at the same store may follow different paths to the backup, gaps may be introduced in the ticket sequence of transactions that are received at the backup stores. This means that (in case of disaster) a transaction that has been properly received may not be able to enter the SUBSCRIBED state because a transaction with a smaller ticket at one store is missing. This implies that T_x will not commit, even if it does not really depend on any

transaction with a smaller ticket. However, one can argue that this disadvantage will not be serious in practice, because the number of transactions unnecessarily aborted will be rather small and limited to transactions that executed just before the catastrophic failure. Given a maximum transmission delay $T_{\max}$ for the transactions to reach the backup site and a maximum processing rate $R_{\max}$, we can bound the number of transactions missed by the backup in case of failure by $T_{\max} \times R_{\max}$.

As we saw, minimum divergence may not be satisfied in the bundle model. The same can happen in the stream model, if the order of messages is not preserved. It is possible to devise a mechanism that strictly enforces the minimum divergence constraint even in these cases, but we believe that its overhead would be too high. To enforce the constraint, each transaction T_i would have to arrive at the backup site with a list of transactions it depends on. Transaction T_i would only be delayed at the backup if the transactions on its list were missing. To construct these lists, the primary site would have to keep track of the last transaction that accessed every lockable database object. This storage and processing cost would be incurred during normal operation. In contrast, our solution pays the price of a few aborted transactions only when a disaster hits. (The minimum divergence constraint could also be satisfied by a mechanism that uses two-phase commit between the two sites.)

Let us now examine our mechanism with respect to the design goals we set in Section 3. Database consistency is preserved as shown in the above observations. The overhead at the primary is minimal: logs are usually maintained for other purposes, so that the only extra processing is the increment of the ticket counter. The mechanism can also scale upwards without limiting performance: there is no component that must process all transactions in some way.

There is not much we can formally prove regarding parallelism, but we can convince ourselves that the proposed mechanism is not a processing bottleneck at the backup site. Transactions are received by the backup hosts in parallel. At each store, locks are requested concurrently by multiple transactions. If the transactions with consecutive ticket numbers at one store access disjoint data sets, then they can acquire their locks at that store in parallel. The only part that is executed serially is the increment of the counter, which is relatively fast. Note that this increment is *not* a global bottleneck, since the increment at one store is independent from the increment at another store, even for the same transaction. On the other hand, if the transactions access common data items at a store, a certain amount of parallelism is lost, because they will acquire their locks sequentially. However, at the primary site, these conflicting transactions also executed sequentially, so the backup is not introducing new delays.

8. INITIALIZATION OF THE BACKUP DATABASE

8.1 Basic Idea

After considering operation under normal conditions, with one site in primary and one site in backup mode, we must now consider the system with

one site in the primary mode and one in the recovering mode. This will be the case at system initialization and when a previously failed site recovers. The mechanism we use is similar to a fuzzy dump [15], which is employed in many commercial systems for media recovery (e.g., Crus [4]). However, there are some differences. In a conventional fuzzy dump one gets a dump of the database and a log of the actions performed while the dump was in progress. In order to restore the database, the dump is installed and the log is replayed against it. No transaction processing takes place while the restoration is performed. The correctness criterion for the restoration process is defined clearly: the database must be brought to the consistent state existing at the time the last transaction recorded on the log committed. Our application is complicated because there are multiple logs and because transaction processing cannot be suspended during recovery: dump generation, dump installation and log playback all occur *simultaneously*. This makes both the implementation and the correctness proof trickier.

The basic idea is for the primary site to scan the entire database and transmit it to the recovering site, along with the changes that occur while this scan is taking place and may therefore not be reflected in the scan. These changes are essentially a redo log and are transmitted over the communication lines used for normal operation. The scan data will probably take too long to be transmitted over the communication lines, so that an alternate path between the primary and the recovering site may be established. This path is usually a tape that is written at the primary site and carried to the backup. The order in which scan messages are received at the recovering site is irrelevant for our method, so that some scan messages may be sent over the communication lines while others are written on tape. Our scheme allows the use of multiple tapes to expedite the process.

8.2 Scanning the Primary

For each primary store S_i, a primary host H_i and a backup host BH_i are selected to copy the S_i database. Host H_i will scan the database at S_i, passing it to BH_i, which in turn installs that data at the corresponding backup store BS_i. First H_i informs S_i that initialization is starting. Then S_i records the current ticket number s_x and returns it to H_i, which in turn sends it (through BH_i) to BS_i. Backup store BS_i sets its ticket number to s_x, creates a legal but empty database and starts accepting redo logs with ticket numbers higher than s_x; redo logs with lower ticket numbers are simply ignored because their effects will be reflected in the scan copy. Store BS_i remains in the recovery mode until initialization completes. If the primary fails while the backup is in recovery mode, nothing can be done: the backup database is still useless.

Under the control of H_i, store S_i starts scanning the portion of the database residing in it. At the same time, normal processing continues and redo logs are sent to the recovering site. The scanning of the database is like a long lived transaction which reads the entire database. For each object scanned (table or record), a scan message is sent to the backup with enough

```
state(T_x) = committed
FOR each table DO
  BEGIN
    get a table read lock;
    send a scan message for table creation;
    FOR each record DO
      BEGIN
        get a read lock on the record;
        send a scan message with the image of the record;
        release the lock held on the record;
      END
    release the lock held on the table;
  END
```

Fig. 5. The scan process.

information for it to create the object. The scan process is described in detail in Figure 5.

Note that the locks are held only while an object is scanned. Records are locked one at a time; groups of records do not have to be locked together. (When the scan process tries to lock an object that is already exclusively locked by a transaction, the scan process does not get blocked: it reads the before image of the record.) Tables or records created by transactions with ticket numbers greater than s_x do not have to be scanned, since they will be transmitted anyway in the redo log of the transaction that created them. However, it will not be harmful if they are transmitted by the scan process, too. The same holds for objects that have been updated after time s_x; it does not matter whether we send the object value that existed at time s_x or at some later time.

8.3 Backup Processing

The backup store will receive data of two types: messages from the scan process and normal redo logs from transactions. The backup store processes both types of messages, using the simple rule of always trying to keep the most recent copy for every object. In particular, when a redo log arrives (with ticket number greater than s_x), it is processed as usual, except for the following:

(1) If a record update action is to be performed but the record does not exist yet, the update is treated as an insert. (This assumes that the update log entries contain the full after image of the modified record.)
(2) If a record (or table) insert arrives, but the record (or table) already exists, the values in the insert replace the existing ones. That is, the end result should be as if the record (or table) did not exist and the insert were normally executed.
(3) Deletes do not actually delete the record (or table). They simply mark it as deleted, but it is still reachable through the primary key index. If a delete arrives, and the object does not exist, a dummy record is inserted with the record key and a flag indicating it is deleted.

When scan messages arrive at BS_i, they are treated as insertions of the record or table. However, if the object already exists (even in deleted form), the scan message is ignored. This is because the scan message may contain obsolete data (i.e., data which have been superseded by a later version). Note that even if the scan message contains a later version than the one already existing, no problem arises, because the later version will be installed when the redo log for the transaction that wrote this later version arrives.

The motivation for the above rules is as follows: during the initialization phase, there will be two types of objects, those accessed by transactions and those not accessed at all. If the object is not accessed, then we want the scan to give us the image of this object as it exists in the primary store. However, if the object is modified by a transaction, then we really do not need the scan for the object since the normal processing will deliver the new image. Our rules ensure that a useless scan does not get in our way. For example, consider a record z that exists at "time" s_x in S_i. The scan starts and sends the image of z. Some time later a transaction deletes z. If the delete arrives at BS_i before the scan, we could end up with a copy of z that should not exist. But since the delete creates a dummy record, the scan message will find it, and the scan copy will be discarded. After initialization, the deleted objects can be removed, so they do not represent a serious storage problem.

8.4 Correctness Arguments

There is a subtle point regarding the time at which we may resume normal processing at the backup site, i.e., exit from the recovery mode. It is not sufficient to wait until the scan process finishes, say at "time" s_y. Let us illustrate this with an example. Suppose that T_i writes data items a and b, which satisfy some consistency constraint (e.g., $a + b = const$). First the scan process holds a read lock on a and transmits its before image (with respect to T_i). Then T_i gets write locks on both a and b, updates them, commits and releases the locks. The scanner gets a read lock on b, transmits its after image and then finishes (say b was the last unscanned object). Suppose that the scan messages for a and b arrive at the backup ahead of the redo log message for T_i. Assume further that the T_i message never makes it because of a failure. If we let the backup site resume normal processing at this point, we end up with an inconsistent copy (we are left with the before image for a and the after image for b).

To avoid this type of problem, we use the following rule. Suppose that the scan process finishes at the primary store S_i when the ticket number is s_y. Then it is safe to resume normal processing at the corresponding backup store BS_i after

(a) all transactions through $T(s_y)$ have committed at the backup store, and
(b) all of the scan messages for this store have been received and processed.

The entire backup site can resume normal processing once all of its stores are ready for normal processing.

Observation: The above two conditions are sufficient for correctness.

Argument: Suppose the scan starts at primary store S_i at "time" s_x (i.e., after transaction $T(s_x)$ ran), and ends at time s_y. Let Z_0 be the state of S_i at time s_x and let T_{scan} be the set of transactions that committed at the primary during the scan.

At the backup store BS_i we start with an empty database state and install a set of transactions T_{back} that includes *all* transactions in T_{scan}. The resulting schedule is SCH_{back}. Concurrently, we process all scan messages and arrive at a database state Z_α. Note that T_{back} may contain more transactions than those in T_{scan}. This is because transactions that committed after time s_y at S_i can arrive before $T(s_y)$.

We want to show that Z_α is identical to Z_β, the state resulting from running SCH_{back} on Z_0. This can be done by considering each object z in the database.

Case I. Object z was not modified by any transaction in T_{back}. Since $T_{scan} \subseteq T_{back}$ (property (a) above), z was not modified by any transaction at the primary site during the scan. Hence, the scan message for z will contain the value of z in Z_0, i.e., $Z_0(z)$. This value will be installed, so $Z_\alpha(z) = Z_0(z)$. This is exactly the value of $Z_\beta(z)$.

Case II. Object z was modified by some transaction in T_{back}. Let T_j be the last transaction in SCH_{back} to have modified z, writing value z_j. That is, $Z_\beta(z) = z_j$. Next, let us look at the moment when z_j is installed by T_j at BS_i. If the scan message for z arrives after this time, it is ignored. If it arrives before this time, T_j overwrites z. In either case, $Z_\alpha(z) = z_j$. Thus, $Z_\alpha(z) = Z_\beta(z)$.

We have shown that after SCH_{back} is run, the state of BS_i is as if we had started with the initial state Z_0. Given the correctness of normal processing, the subsequent states of BS_i will also have this property. This means that after conditions (a) and (b) hold, our implicit assumption about initial states holds (see Section 4).

9. COPING WITH SINGLE HOST OR STORE FAILURES

We have only addressed disasters so far. However, partial failures are also possible, and one would expect them to occur much more frequently than disasters, so it would be desirable for the remote backup to help the primary recover from such failures as well. In this section we consider some failure scenarios for single components and discuss what can be done in each case.

Let us consider the failure of a primary host H_i. As we mentioned in Section 3, the architectural configuration plays an important role in the failure scenario. We first examine the case where stores and hosts are different processors and each host has independent access to each store. Under this assumption, the failure of the host has no impact on the data; only processing power is lost. The transactions that would be processed by H_i will now be spread among the remaining hosts to achieve graceful performance degradation. It is also necessary to take care of the transactions that were being processed by H_i and were still in progress when the failure occurred. A

replacement host RH_i is selected from the remaining hosts at the primary site to perform this task. Host RH_i is in charge of committing or aborting pending H_i transactions. In addition, if the stream model is used for log propagation, RH_i will start acting as the supervising host for stores which H_i had supervised. Finally, RH_i is responsible for sending to the backup any redo logs that should have been sent by H_i but did not arrive because of H_i's failure. In order to find such logs, RH_i contacts the stores and obtains the logs for committed transactions that should have been propagated by H_i but have not been acknowledged by the backup (see Section 5). As in normal processing, when an acknowledgement is received that a transaction has successfully committed at the backup site, RH_i informs the primary stores to purge the corresponding log entries.

Similarly, if a backup host BH_i fails, a replacement host RBH_i is selected. Host RBH_i contacts the backup stores and obtains the status of all transactions in progress that were being coordinated by BH_i; RBH_i becomes the coordinator for these transactions and tries to commit them. Host RBH_i is also responsible for obtaining the logs that were sent to BH_i while it was down and before RBH_i took over. This is done by contacting the appropriate primary host(s) and obtaining the redo logs that have not been acknowledged by the backup. Again, in the stream model, host RBH_i must act as the supervising host for the stores that BH_i had supervised.

In the previous sections we assumed that the stores were fault tolerant. If this is not true, a store failure at the primary may render part of the data inaccessible. This is also the case if a failure occurs in the model where stores and hosts are not disjoint processors, but a computer acts partly as a host and partly as a store. It is desirable to handle such a failure without declaring a disaster. One possibility would be for the backup store corresponding to the failed primary store to be "promoted" to primary and join the rest of the primary stores to form a new primary processing group. Our algorithm does not handle promotions and substantial changes would have to be made. We do not address those here; instead, we simply note that the following issues must be considered by a promotion algorithm:

—First, special care must be taken to preserve the consistency of the database. Some transactions that committed at the failed store may not have reached the backup. When the new primary processing group is formed, the changes made by some of these transactions will not be reflected in the promoted store but will appear in the rest of the stores. Thus, transaction atomicity may be violated.

—Second, processing at the backup must be modified. It is now necessary to handle transactions that access data at the promoted store in a special way, since that store is no longer available for backup processing.

10. EXTENSIONS

In the preceding sections we assumed that two-phase locking was used for concurrency control. Two-phase locking is conceptually simple and has been studied extensively, so that adopting it in our model made the presentation

simpler and helped us concentrate on the novel issues. However, two-phase locking may not be the method of choice for some real systems (e.g., O'Neil [14]). We now show that our algorithm can also be applied to systems using other concurrency control mechanisms.

The only requirement that must actually be satisfied to make our backup method work is that the concurrency control algorithm assign tickets to transactions at each store such that they have the following property: if at the primary two *logical* actions A_1 and A_2, executed by transactions T_x and T_y respectively, induce a dependency $T_x \rightarrow T_y$, then $ticket(T_x) < ticket(T_y)$. As long as this condition holds, one can see that our proofs are valid, no matter what the details of the particular concurrency control algorithm are. Logical actions are record reads, writes, etc., as discussed in Section 3. We would like to emphasize that we are concerned with logical actions *only*; we do not care about actions at the physical level (e.g., compacting records within a disk page).

The above requirement implies that the schedule of logical actions must be serializable. For every concurrency control method that generates serializable schedules there is a way to figure out a serial order (this follows from the fact that serializable schedules can be topologically sorted [2]). In two-phase locking this can be done easily by incrementing a counter before starting to release locks at commit time. In timestamp ordering, the timestamps themselves can be used as tickets.

In other concurrency control methods it may be more difficult to determine (on-line) an equivalent serial order for transactions and use that as ticket. In such cases, one may use multiple tickets for a transaction (each for a set of data accessed by a transaction). These tickets have a limited scope, i.e., they are only used at the backup to determine the order of conflicting accesses to the set of data items for which the ticket was issued at the primary. The particular way to assign such tickets depends on the details of the concurrency control method and is not discussed here.

A different extension is to break up the data residing in a store into *chunks* and apply our proposed solution, generating ticket numbers per chunk instead of per store. For example, each relation could be a chunk. This has the advantage of gaining parallelism, by reducing the critical sections: there is less contention for getting the ticket numbers at the primary stores, less contention for the state flipping (from LOCKING to SUBSCRIBED) at the backup stores, etc. In addition to this, the impact of a missing transaction in case of failure is reduced. Let us illustrate with an example. Suppose disaster strikes and transaction T_x does not make it to the backup site. All transactions accessing data items in the same chunks as T_x that have been received properly at the backup but have higher ticket numbers than T_x have to be aborted. The less data that the chunks accessed by T_x contain, the fewer aborted transactions. On the other hand, more ticket numbers have to be processed both at the primary and the backup site, which will add some overhead.

Another optimization can be made in the scanning process. As we saw in Section 8, some of the messages sent by the scanning process are ignored at

the backup site. We can decrease the number of such messages if we keep track of what has been scanned at S_i. For simplicity assume that each record has a scan bit that indicates if it has been scanned. (This might actually be implemented with a cursor showing how far the scan process has gone.) As each record is scanned, the bit is set. If a transaction modifies a record that has not been scanned, the after image for that record need not be sent, since the scan process will send the after image of the record later. However, the redo log entry containing the ticket must be sent, so that the ticket sequence will not be broken.

Yet another interesting problem is the partition of the data among the stores. In this paper we took the partition to be identical at both sites. This is the most natural case (since the backup system will probably be a replica of the primary). However, it turns out that arbitrary partitions are not possible, at least within our framework. Consider the following case: backup store BS_{ij} contains data from primary stores S_i and S_j. Suppose that $ticket(T_x) < ticket(T_y)$ at S_i and $ticket(T_x) > ticket(T_y)$ at S_j (at one of the primary stores there is no dependency between the two transactions). It is possible that all of the data accessed by T_x and T_y resides on BS_{ij}. Which ticket number will be used to determine the order in which the transactions will enter the SUBSCRIBED state at BS_{ij}? This leaves only the possibility of a *finer* partition at the backup site. But it turns out that even this is not practical. Suppose the data on primary store S_i is split into backup stores BS_{i1} and BS_{i2} at the backup. Transaction T_x accesses data only on BS_{i1} and has $ticket(T_x) = k$. Transaction T_y accesses data only on BS_{i2} and has $ticket(T_y) = k + 1$. Transaction T_x will never appear in store BS_{i2}. This means that transaction T_y will never execute, because it cannot enter the SUBSCRIBED state unless T_x has previously done so! Thus, the partitions of data at the two sites have to be identical. Note that the partitions only have to be *logically* identical. In other words, for every process assigning tickets at the primary, there must be a peer process at the backup counting the tickets and controlling the same data as the process at the primary. The physical distribution of the data among computers and the physical layout on storage media can be different, as long as the same data at the two sites is controlled by corresponding processes. Furthermore, no restriction applies to the way the data is partitioned: arbitrary partitions are allowed, as long as the partition at the two sites is the same.

The method for recovery we suggested in Section 8 assumes that the primary database is lost during a disaster. If this is not the case, it may be possible to bring the failed primary up-to-date by sending it the redo logs for the transactions it missed, instead of a copy of the entire database. The recovering primary also has to undo transactions that did not commit at the backup before redoing the missed transactions. The use of redo logs for recovery may or may not be feasible, depending on the time it takes for the failed site to be repaired. If this time is long, the logs will grow too big. The relative performance of the two methods also depends on the size of the logs. Our method is general and works even when the primary loses its copy entirely, without excluding the use of the other method. For example, one

could combine the two strategies as follows: when a failure of the primary is detected, the backup takes over transaction processing and tries to keep the logs for as long as it can, in case the primary recovers soon and still has its copy. When its capacity overflows, it starts discarding the logs; when the failed site recovers, the method proposed in Section 8 will be used.

11. CONCLUSIONS

We presented a method for keeping a remote backup database up-to-date for disaster recovery. The method ensures that the backup copy will be consistent with the primary and that in case of failure the backup copy will be (at most) a few transactions behind the primary. The method is relatively straightforward and can be implemented using well known concepts and techniques, such as locking and logging. The overhead imposed at the primary site is relatively small, and there is *no central processing* in our mechanism, i.e., no component that must "see" all transactions. This means that the system can scale upwards: more communication lines, hosts and stores can be added without having backup management interfere.

ACKNOWLEDGMENTS

The authors would like to thank Jim Gray, Robert Hagmann, C. Mohan, Andreas Reuter and the referees for their comments and the references they provided.

REFERENCES

1. AGRAWAL, R. A parallel logging algorithm for multiprocessor database machines. In *Proceedings of the 4th International Workshop on Database Machines*. Springer, New York, 1985.
2. BERNSTEIN, P. A., HADZILACOS, V., AND GOODMAN, N. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley, Reading, Mass., 1987.
3. BURKES, D., AND TREIBER, K. Design approaches for real time recovery. Presentation at the *Third International Workshop on High Performance Transaction Systems* (Pacific Grove, Calif., Sept. 1989).
4. CRUS, R. A. Data recovery in IBM Database 2. *IBM Syst. J. 23*, 2 (1984), 178–188.
5. FINKELSTEIN, W., AND CAPPI, M. Experiences with large networks of computers. Presentation at the *International Workshop on High Performance Transaction Systems* (Pacific Grove, Calif., Sept. 1985).
6. GARCIA-MOLINA, H., AND ABBOTT, R. K. Reliable distributed database management. In *Proceedings of the IEEE, Special Issue on Distributed Database Systems* (May 1987), 601–620.
7. GRAY, J. N., AND ANDERTON, M. Distributed computer systems: Four case studies. In *Proceedings of the IEEE, Special Issue on Distributed Database Systems* (May 1987), 719–726.
8. GRAY, J. N. Why do computers stop and what can be done about it? Presentation at the *Fifth Symposium on Reliability in Distributed Software and Database Systems* (Los Angeles, Calif., Jan. 1986).
9. GRAY, J. N. Notes on database operating systems. *Operating Systems: An Advanced Course*. R. Bayer et al., Eds., Springer, New York, 1979.
10. GRAY, J. N., AND REUTER, A. Transaction processing. *Course Notes from CS # 445 Stanford Spring Term*, 1988.
11. KORTH H. F., AND SILBERSCHATZ, A. *Database System Concepts*. McGraw-Hill, New York, 1986.

12. *IBM, IMS/VS Extended Recovery Facility (XRF)· General Information.* Doc. GG24-3150, March 1987
13. LYON, J Design considerations in replicated database systems for disaster protection. *IEEE Compcon*, 1988
14. O'NEIL, P. E. The escrow transactional method. *ACM, Trans Database Syst 11*, 4 (Dec. 1986), 405-430.
15. ROSENKRANTZ, D. J. Dynamic database dumping. In *Proceedings of SIGMOD International Conference on Management of Data.* ACM (1978), 3-8
16. SKEEN, D. Nonblocking commit protocols In *Proceedings of the ACM SIGMOD Conference on Management of Data* (Orlando, Fl., June 1982), 133-147
17. SCHLICHTING, R. D , AND SCHNEIDER, F D. Fail-stop processors: An approach to designing fault-tolerant computing systems *ACM, Trans. Comput Syst 1* (Aug 1983), 222-238.
18. Tandem Computers *Remote Duplicate Database Facility (RDF) System Management Manual.* March 1987
19. TANENBAUM, A. S. *Computer Networks.* Prentice Hall, Englewood Cliffs, N.J , 1988.

Received October 1989; revised March 1990; accepted April 1990