

# MUREX: A Mutable Replica Control Scheme for Structured Peer-to-Peer Storage Systems

Jehn-Ruey Jiang

*Department of Computer Science and Information Engineering, National Central University, Zhongli, 320 Taiwan*  
Chung-Ta King, Chi-Shiang Liao and Ching-Hao Liu

*Department of Computer Science, National Tsing Hua University, Hsinchu, 300 Taiwan*

*Received: —; revised: —*

**Abstract**—This paper proposes MUREX, a mutable replica control scheme, to keep one-copy equivalence for synchronous replication in structured P2P storage systems. For synchronous replication in P2P networks, it is proper to adopt crash-recovery as the fault model; that is, nodes are fail-stop and can recover and rejoin the system after synchronizing their state with other active nodes. In addition to the state synchronization problem, we identify other two problems to solve for synchronous replication in P2P storage systems. They are the replica acquisition and the replica migration problems. On the basis of multi-column read/write quorums, MUREX conquers the problems by the replica pointer, the on-demand replica regeneration, and the leased lock techniques. We prove the correctness of MUREX, analyze and also simulate it in terms of communication cost and operation success rate.

## I. INTRODUCTION

Peer-to-peer (P2P) systems are based on self-organizing, decentralized overlay networks, in which participating peer nodes play symmetric roles — both servers and clients. P2P systems are usually designed to accommodate a large number of nodes and to adapt to dynamic node joining and leaving. The most well-known application of P2P systems is the storage data sharing over Internet, as typified in P2P music file sharing [6], [17]. However, file sharing is only one of the functions that a storage system can support. Following the great success of P2P file sharing, a natural next step is to develop wide-area, P2P storage systems to aggregate idle storage across the Internet to be a huge storage space.

P2P storage system has been an active research topic and many systems have been proposed [2], [7]–[9], [11], [16], [19], [21], [24]. Some systems adopt the *unstructured P2P* approach [2], [7], in which there is no restriction on the interconnection of the nodes. Unstructured P2P storage systems are easy to build and maintain, but it is difficult to guarantee the quality in accessing the stored data [18]. Many P2P storage systems [8], [9], [11], [16], [19], [21], [24] are thus built on top of *structured P2P networks* [18], [20], [22], [25].

Structured P2P storage systems rely on a hashing scheme (viz., a hash function mapping) to name the peer nodes. With the hashed ID, the peer nodes can be linked with a certain interconnection structure, such as a ring, a torus, or a hypercube. The data objects to be stored in the storage system

are also named with the same hash function. A data object with the hashed key  $k$  is published to and managed by the peer node whose hashed key (or ID) is “closest” to  $k$ . To fetch a data object with hashed key  $k$ , a request is routed according to the interconnection structure until the node with the closest hashed ID is reached. In this way, any data object in the storage system can be located and accessed within a certain bound of message relays, no matter where the request is initiated. This in essence builds a *distributed hash table* (DHT) across the participating nodes.

With the distributed hash table, structured P2P systems can handle dynamic node joining and leaving. Note that according to the data allocation scheme described above, any given key in the hash table will have a node taking charge of the corresponding entry. Even after that node leaves, the underlying routing scheme will always send requests for that key to the node currently having the closest ID to the key. In this way, the leaving node is substituted and the keys managed by it are taken over by the substituting node. Please refer to Fig. 1 for such a scenario (node  $p$  substitutes leaving node  $q$ ). Similarly, when a node newly joins the network, it will partially substitute a certain node to manage the keys that are now closest to its ID. Please refer to Fig. 1 for such a scenario (newly joining node  $u$  partially substitutes node  $v$ ).

Although the underlying P2P routing can adapt to dynamic node joining and leaving, there is a problem that the data object stored in nodes will be lost when nodes fail or leave. A common solution to this problem is to replicate the data objects among nodes to provide high data availability. If the data objects are read-only (or non-mutable), then the P2P storage system will only need to consider where to replicate the data objects [8], [9], [11]. The system becomes much complicated if the data objects are mutable [16], [19], [21], [24]. In this paper, we concentrate on mutable P2P storage systems because they are desirable by most practical applications.

In mutable P2P storage systems, data replication must obey the criteria of one-copy equivalence to ensure data consistency. There are two types of mechanisms to achieve such a criterion: synchronous replication and asynchronous replication. The former requires that each write operation should finish updating all replicas before the next write operation proceeds. The latter regards a local write operation as

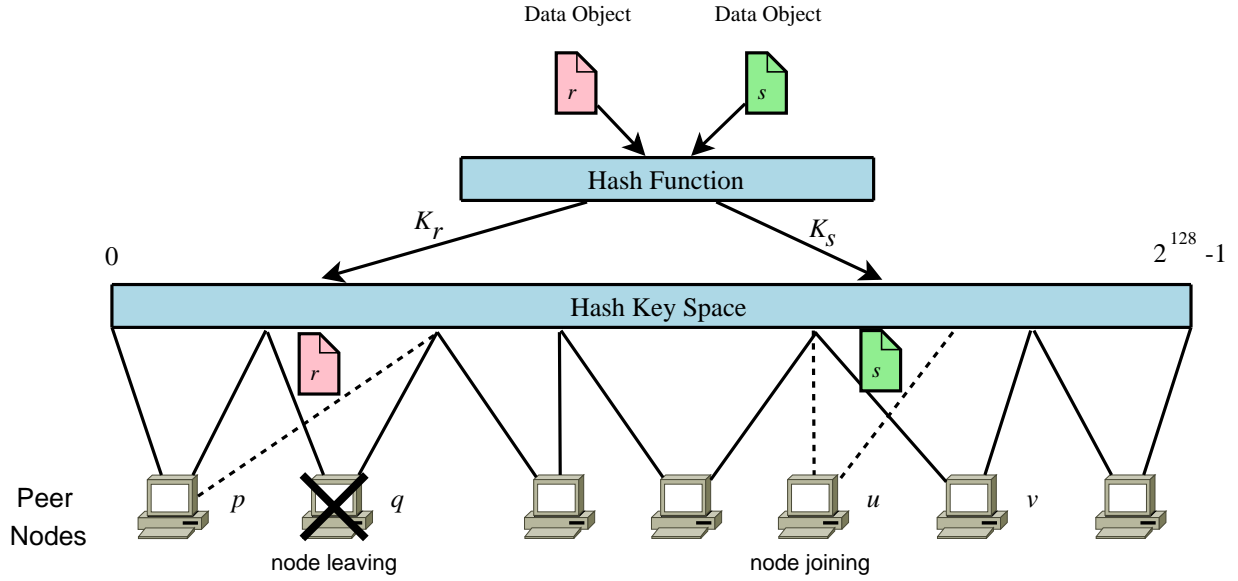


Fig. 1. The scenarios of node joining and leaving.

complete once data object is written to the local replica; data object is then asynchronously written to other replicas. The synchronous replication can ensure data consistency strictly, but may have long operation latency. On the other hand, the asynchronous replication may violate data consistency, but has shorter latency. However, when data inconsistency occurs, complex log-based mechanisms should be invoked to roll back the system to a consistent state. In this paper, we adopt synchronous replication since we take data consistency as the most significant factor and we regard that there may not be available storage for storing logs for system roll-back in asynchronous replication.

For synchronous replication in P2P networks, it is proper to adopt crash-recovery as the fault model. In a crash-recovery system, nodes are fail-stop and can recover and rejoin the system after synchronizing their state with other active nodes. In addition to the *state synchronization problem*, we have two more problems to solve for synchronous replication in P2P storage systems. First, in P2P environments, an active node  $p$  may substitute some failing/leaving node  $q$  in the recovery process. Thus, node  $p$  must acquire the replicas hosted by node  $q$  somehow. Below, we call this the *replica acquisition problem*. Second, a newly joining node  $u$  will partially substitute an active node  $v$  to share  $v$ 's load by hosting part of  $v$ 's replicas. Thus, part of  $v$ 's replicas should be migrated to node  $u$ . Below, we call this the *replica migration problem*.

In this paper, we propose MUREX, a mutable replica control scheme, to keep one-copy equivalence for synchronous replication in structured P2P storage systems. On the basis of multi-column read/write quorums, MUREX conquers the problems mentioned by the *replica pointer*, the *on-demand replica regeneration*, and the *leased lock* techniques. We will prove the correctness of MUREX, analyze and also simulate it in terms of communication cost and operation success rate.

The rest of the paper is organized as follows. Preliminaries

are given in section II. In section III, we discuss the problems encountered in realizing synchronous replication for P2P mutable storage systems. We then in section IV show how MUREX solves the problems, and analyze and simulate MUREX in section V. Some related works are introduced in section VI, and concluding remarks are drawn in section VII.

## II. PRELIMINARIES

As a replica control scheme, MUREX needs to ensure data consistency. In this paper, we adopt the *one-copy equivalence* consistency criteria, which states that the set of replicas must behave as if there were only a single copy. Conditions to ensure one-copy equivalence are

- 1) no pair of write operations can proceed at the same time,
- 2) no pair of a read operation and a write operation can proceed at the same time,
- 3) a read operation always returns the replica that the last write operation writes.

*Quorum-based schemes* are popular mechanisms to enforce one-copy equivalence for synchronous replication since they render relatively high data availability and low communication cost. The basic concept of such schemes is described as follows. Each data object has  $n$  replicas, each of which is associated with a version number. A read operation should read-lock and access a read quorum of replicas and return the replica owning the largest version number. On the other hand, a write operation should write-lock and access a write quorum of replicas and then updates them with the new version number, which is one more than the largest version number just encountered. If we restrict the write-write exclusion and the write-read lock exclusion, and restrict that any pair of a read quorum and a write quorum, and any two write quorums have a non-empty intersection, then one-copy equivalence is guaranteed.

There are several mechanisms proposed in the literature for forming read and write quorums, such as tree quorums [1],

majority quorums [10], [23], grid quorums [4], and multi-column quorums [13], etc. MUREX adopts the multi-column quorums, which have the smallest quorums (constant-sized quorums in the best case) among the mechanisms. It is noted that smaller quorums imply few accesses of replicas, which in turn imply lower communication cost. Furthermore, as shown in [13], multi-column quorums are candidates to achieve the highest availability, which is the probability for a quorum to be formed in an error-prone environment.

Multi-column quorums are constructed with the aid of the multi-column structure  $MC(m) \equiv (C_1, \dots, C_m)$ , which is a list of pairwise disjoint sets of replicas. Each set  $C_i$  is called a *column* and must satisfy  $|C_i| > 1$  for  $1 \leq i \leq m$ . For example,  $(\{r_1, r_2\}, \{r_3, r_4, r_5\}, \{r_6, r_7, r_8, r_9\})$  and  $(\{r_1, r_2, r_3, r_4, r_5\}, \{r_6, r_7\}, \{r_8, r_9\})$  are multi-column structures, where  $r_1, \dots, r_9$  are replicas of a data object.

By organizing data replicas as multi-column structure  $MC(m) \equiv (C_1, \dots, C_m)$ , the write and the read quorums are defined as follows:

A *write quorum* under  $MC(m)$  is a set that contains all replicas of some column  $C_i$ ,  $1 \leq i \leq m$  (note that  $i = 1$  is included), and one replica of each of the columns  $C_{i+1}, \dots, C_m$ .

A *read quorum* under  $MC(m)$  is either

Type-1: a set that contains one replica of each of the columns  $C_1, \dots, C_m$ , or

Type-2: a set that contains all replicas of some column  $C_i$ ,  $1 < i \leq m$  (note that  $i = 1$  is excluded), and one replica of each of the columns  $C_{i+1}, \dots, C_m$ .

For example, under  $MC(2) = (\{r_1, r_2, r_3\}, \{r_4, r_5\})$ , the possible write quorums are  $\{r_4, r_5\}$ ,  $\{r_1, r_2, r_3, r_4\}$ ,  $\{r_1, r_2, r_3, r_5\}$ , and the possible read quorums are  $\{r_1, r_4\}$ ,  $\{r_1, r_5\}$ ,  $\{r_2, r_4\}$ ,  $\{r_2, r_5\}$ ,  $\{r_3, r_4\}$ ,  $\{r_3, r_5\}$  (of type-1) and  $\{r_4, r_5\}$  (of type-2). Note that the write quorum definition and the type-2 read quorum definition are identical except that the latter does not include the sets composed of all replicas in  $C_1$  and one replica from each of  $C_2, \dots, C_m$ .

As we have mentioned, multi-column quorums have constant size in the best case. Below, we discuss the sizes for quorums under the multi-column structure  $MC(m)$  whose columns all have the size  $s$ , where  $m > s$ . It is noted that we will denote such a multi-column structure as  $MC(m, s)$  in the following context. The write quorum under  $MC(m, s)$  has constant size  $s$  in the best case and  $s + m - 1$  in the worst case. The read quorum under  $MC(m, s)$  has constant size  $s$  in the best case and size  $s + m - 2$  in the worst case.

### III. THE PROBLEMS

In this section, we identify three problems encountered in enforcing synchronous replication for structured P2P storage systems. The three problems are replica migration, replica acquisition, and state synchronization. Below, we elaborate the problems one by one.

- *Replica Migration*: When a node  $u$  newly joins the system and partially substitutes another node  $v$  to host some replicas, node  $v$  should transfer the replicas to  $u$  immediately. However, in a constantly changing P2P

environment, the cost of transferring replicas may be too high. We need an efficient mechanism to allow replicas to migrate from substituted node to substituting node.

- *Replica Acquisition*: When an active node  $p$  substitutes a failing/leaving node  $q$ , node  $p$  needs to acquire all replicas hosted by  $q$ . The problem is that node  $q$  has no idea about which replicas are hosted by  $p$ . Thus, we need a mechanism to make node  $p$  know which replicas are hosted by node  $q$  and to acquire the replicas efficiently.
- *State Synchronization*: Suppose an active node  $p$  substitutes a failing/leaving node  $q$ , and  $p$  has acquired a replica  $r$  hosted by  $q$  previously. To make replica  $r$  effective, we have to synchronize  $r$ 's state, i.e., to ensure that all the participating nodes have the same view with respect to  $r$ 's states. We must ensure the acquired replica  $r$  is an up-to-date copy. Furthermore, since there may be a node that has locked replica  $r$  to make  $r$  in the locked state, we need a mechanism to ensure that the locked state is not violated after  $p$  acquires replica  $r$ .

### IV. THE PROPOSED SCHEME AND THE CORRECTNESS

In this section, we introduce the proposed scheme — MUREX, which uses the read/write quorum, the replica pointer, the on-demand replica regeneration, and the leased lock techniques to solve the three problems mentioned in the last section. We also show in this section that MUREX can ensure the one-copy-equivalence criterion for synchronous replication in P2P storage systems. Below, we first give an overview of MUREX and then elaborate its details.

#### A. Overview

For a data object, there are  $n$  replicas with hashed keys  $k_1, \dots, k_n$ , where  $k_1 = \text{HASH}_1(\text{data object name}), \dots, k_n = \text{HASH}_n(\text{data object name})$ . The replicas are disseminated to the nodes whose hashed ID are nearest to  $k_1, \dots, k_n$ , respectively. Please refer to Fig. 2 for the illustration of the replica dissemination. Each replica has a version number which is 0 initially and will increase gradually. MUREX organizes the  $n$  replicas into a multi-column structure to help form read and write quorums. A quorum is a subset of the nodes storing the  $n$  replicas; it should satisfy the *intersection property* that any pair of a read quorum and a write quorum, and any two write quorums have at least one common member. MUREX provides the following operations:

- *publish(CON, DON)*: to place  $n$  replicas at the nodes associated with  $k_1, \dots, k_n$  for the object of name DON (standing for Data Object Name) with content CON (standing for CONtent) and version number 0.
- *read(DON)*: to acquire the up-to-date replica of the object of name DON by locking all replicas of a read quorum.
- *write(CON, DON)*: to update all the replicas of a write quorum with content CON for the object of name DON.

#### B. Read/Write Quorum Construction

Initially, a data object owner publishes the data object of name DON with content CON by calling *publish(CON,*

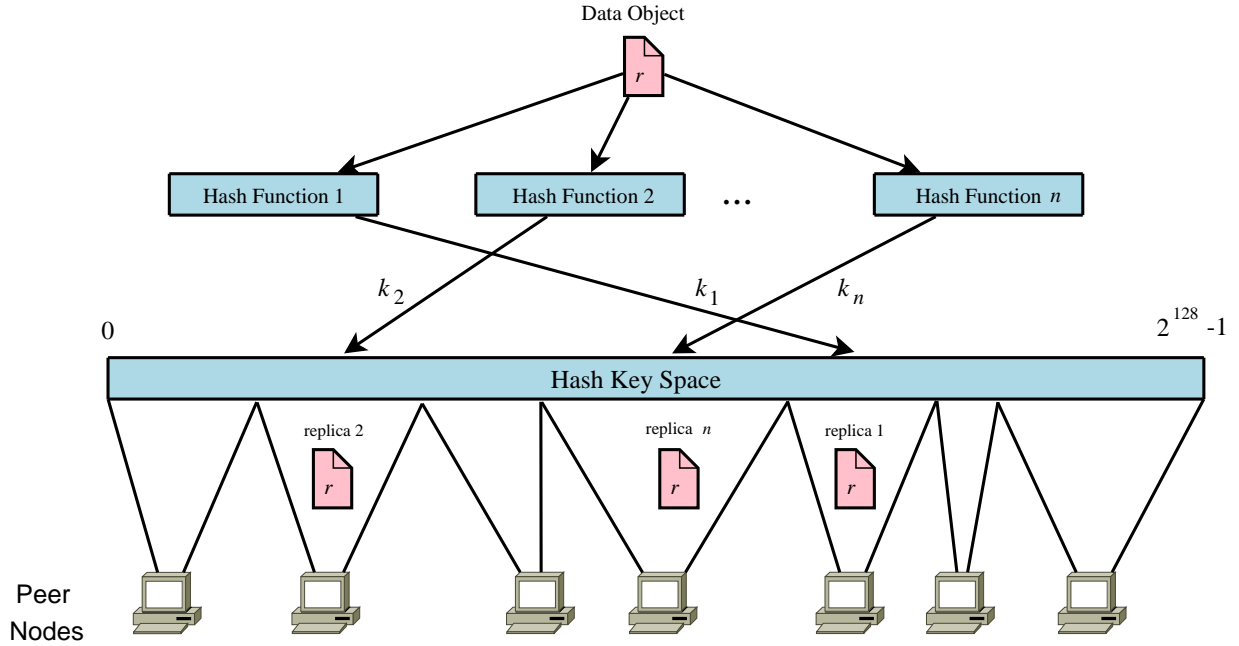


Fig. 2. The dissemination of  $n$  replicas of a data object.

DON). Afterwards, any participant can call *read* (or *write*) operation to read (or write) the data object by issuing RLOCK (or WLOCK) messages. With the help of the DHT and the multi-column structure, the messages will reach all members of a read (or write) quorum. The two functions *Get\_Read\_Quorum* and *Get\_Write\_Quorum* in Fig. 3 try to issue RLOCK and WLOCK messages to nodes in a last-to-first column-wise manner to return respectively the read and the write quorums under a multi-column structure. It is noted that below we use the words “node” and “replica” interchangeably since a replica must be hosted by a node. Below, we also use LOCK message to stand for WLOCK message or RLOCK message.

When a node receives a LOCK message to request for locking a data object, it sends a MISS message to the requester if it does not own a replica of the data object. As we will show, the MISS message will cause the requester to send an up-to-date replica of the data object later. It is noted that a node has at most one pending MISS message for each replica. A MISS message is said to be pending if there is no replica sent in response to it. When a node has a pending MISS message for a replica and further receives LOCK message for locking the replica, it will send WAIT message to the requester. On the other hand, if the node owns the replica when it receives a LOCK message for a data object, it then checks if there is a lock conflict. We say that there is a lock conflict if a read-locked replica receives a write-lock request, or if a write-locked replica receives a write-lock or a read-lock request. If there is no lock conflict, the node locks the replica and then replies with OK message containing the replica version number. On the contrary, if there is a lock conflict, the node replies with WAIT message.

After sending LOCK messages, a node enters the “wait period”, which is of the length of a turn-around time. During

the wait period, if a node has got any WAIT message, it can conclude that there is lock contention. For such a case, the node sends UNLOCK messages to all the nodes that it has sent LOCK messages. Only after a random backoff time, can the node start over again to send LOCK messages for locking replicas of a quorum. The random backoff concept is borrowed from Either Net [5]. It is used to avoid continuous conflicts among contending nodes.

After *Get\_Write\_Quorum* or *Get\_Read\_Quorum* function returns a write quorum or a read quorum, it means that all replicas in the quorum have been locked. The node calling the function can then execute the desired operation. After the operation is finished, a node sends UNLOCK messages to all nodes that it has sent LOCK messages to unlock the replicas. A read operation in MUREX reads the replica of the largest version number from the read quorum. On the other hand, a write operation always writes all replicas of a write quorum with the version number one more than those ever encountered.

### C. Replica Pointers

When a node  $u$  newly join the system to share part of the load of node  $v$  by managing replicas of keys from  $k$  to  $k'$ , the replicas of keys from  $k$  to  $k'$  should migrate from  $v$  to  $u$ . To reduce the cost of transferring all the replicas, MUREX transfers replica pointers instead of the actual replicas. A *replica pointer* is a five-tuple of the form:  $\langle \text{hashed key, data object name, version number, lock state, actual storing location} \rangle$ . It is produced when a replica is generated and can be used to locate the actual replica. When node  $v$  owns the replica pointer of replica  $r$ , it is regarded as  $r$ 's host, which can reply to the lock request for  $r$ . On the other hand, when node  $v$  sends out the replica pointer of replica  $r$ , it is no more

```

function GET_WRITE_QUORUM((( $C_1, \dots, C_m$ ): MC Structure): Set;)
  Var  $R = \emptyset$ : Set;
  for  $i = m, \dots, 1$  do
    Send WLOCK to all nodes in  $C_i$  and enter “wait period” for getting replies;
    if all nodes in  $C_i$  reply with WAIT or MISS then
      Send UNLOCK to nodes in  $C_1 \cup \dots \cup C_i$ ; Exit;
    end if
    if all nodes in  $C_i$  reply with OK then
      return  $R \cup C_i$ ;
    else if a node  $u$  replies OK then
       $R = R \cup \{u\}$ ; //continue the for loop here
    end if
  end for
end function

function GET_READ_QUORUM((( $C_1, \dots, C_m$ ): MC Structure): Set;)
  Var  $R = \emptyset$ : Set;
  for  $i = m, \dots, 1$  do
    Send RLOCK to all nodes in  $C_i$  and enter “wait period” for getting replies;
    if all nodes in  $C_i$  reply with WAIT or MISS then
      Send UNLOCK to nodes in  $C_1 \cup \dots \cup C_i$ ; Exit;
    end if
    if  $i \neq 1$  and all nodes in  $C_i$  reply with OK then
      return  $R \cup C_i$ ;
    else if  $i \neq 1$  and a node  $u$  replies with OK then
       $R = R \cup \{u\}$ ; //continue the for loop here
    else if  $i = 1$  and a node  $u$  replies with OK then
      return  $R \cup \{u\}$ ;
    end if
  end for
end function

```

Fig. 3. Two functions that can properly return a read and a write quorum, respectively.

the host of  $r$  and cannot reply to the lock request for  $r$  (even if it stores the actual replica of  $r$ ).

The replica pointer is a lightweight mechanism for transferring replicas; it can be propagated from node to node in a very low cost. When a node  $u$  owning the replica pointer of  $r$  receives a lock request for  $r$ , it should check whether the node actually storing  $r$  is still alive. If so,  $u$  can behave as host of  $r$ . Otherwise,  $u$  regards itself as having no replica  $r$ . It is noted that every transfer of replica pointer between two nodes, say from  $v$  to  $u$ , should be recorded locally by  $v$  so that later messages, such as UNLOCK message, destined to  $v$  for replica  $r$  can be sent to the last node having the replica pointer.

#### D. On-Demand Replica Regeneration

When a node  $q$  fails/leaves and another node  $p$  substitutes node  $q$ , it is needed for node  $p$  to acquire all replicas hosted by  $q$ . However, we have the problem that node  $p$  has no idea about which replicas are hosted by  $q$ . Below, we show how the replicas can be acquired in an on-demand manner. The term “on-demand” means that node  $p$  only acquire requested replicas. When node  $p$  receives from node  $u$  LOCK message for locking a replica, it should send MISS message if it does

not own the replica. Node  $p$  is assumed to have no replica  $r$  if the following conditions hold:

- 1)  $p$  does not have the replica pointer of  $r$ .
- 2)  $p$  has the replica pointer of  $r$  and the pointer indicates that  $w$  stores  $r$ , but  $w$  is not alive.

After obtaining (resp., generating) the newest replica by executing a read (resp., write) operation, node  $u$  should send the replica to node  $p$ . It is noted that a node has at most one pending MISS message for a replica. Furthermore, when a node has a pending MISS message for a replica and further receives a lock request for the replica, it will send WAIT message to the requester. In such a manner, we can ensure that a node will only receive one replica in response to MISS message.

By the on-demand replica regeneration technique, node  $p$  passively acquires replicas only when the replicas are requested. For the replicas never requested, there is no need to acquire them to keep the overhead as low as possible. However, the number of replicas of a data object may decrease gradually and influence the persistency of the data object. Fortunately, the bad influence does not occur for replicas that are accessed frequently. Moreover, we can allow the publisher of a data object to periodically perform the “dummy read

operation” for the data object, which will be described later. We even can demand each participating node to periodically perform the dummy read operation for rarely-accessed data object replica hosted by it. When a replica of a data object is not accessed for a specific period of time, the dummy read operation is performed once. The dummy read operation is similar to the read operation and plays the role of checking if replicas of the data object are still alive; it does not read the replica in practice and thus only incurs little overhead. When some replicas of the data object are missed, the node initiating the dummy read operation can re-disseminate the replica to the proper node. The persistency of the data object can thus be ensured.

### E. Leased Locks

When a replica  $r$  of a data object is re-disseminated to some node, we must ensure that all participating nodes have the same view with respect to the replica. We first need to ensure the replica is up-to-date. If the replica is re-disseminated due to a node’s receiving MISS message, the replica is surely up-to-date. This is because a node re-disseminates the replica only after it has executed the read (or write) operation to acquire (or generate) the up-to-date replica. On the other hand, if the replica is re-disseminated due to a node’s performing a dummy read operation, the node is demanded to first obtain the up-to-date replica and then to re-disseminate the replica.

The second thing for all participating nodes to agree with is the state of replica  $r$ . Since there may be some node that has locked replica  $r$  to make  $r$  in the lock state, we need to ensure that the lock state is not violated. To achieve this, each lock is assumed to be a *leased lock* that has a leased period of  $L$ . That is to say, after a replica is locked, it becomes unlocked automatically after a period of  $L$ . Assume that the critical section (CS) of a read or a write operation takes  $C$  time to complete. A node should release any obtained lock if it still has no chance to enter the CS and  $H > L - C - D$  holds, where  $H$  is the holding time of the lock and  $D$  is the propagation delay for transmitting the lock. Please see Fig. 4 for the relationship among  $H$ ,  $L$ ,  $C$  and  $D$ . The condition of  $H > L - C - D$  can ensure a node to complete the desired operation before any lock expires.

When a node detects that a lock of a specific replica is expiring (i.e.,  $H > L - C - D$  is going to hold), it is possible that the locks of other replicas will also expire in the near future. Thus, we demand a node to release all locks and start over to acquire the locks again. In this manner, MUREX can avoid deadlock and starvation. Furthermore, we demand a node to wait for a random backoff time before acquiring the locks next time. This can alleviate the chance of repeatedly occurrence of contention-then-release-all-locks situation.

Now, we describe how to make all participating node have the same view for the lock state by the concept of leased locks. Suppose a node  $p$  substitutes a failing/leaving node  $q$  to host replica  $r$ , and node  $p$  has received the up-to-date replica of  $r$  at time  $T$ . After receiving the up-to-date replica, node  $p$  generates a replica pointer for  $r$  and can start to reply to LOCK message for locking  $r$  at time  $T + L$ , where  $L$  is the leased period of

the lock. In this manner, all participating nodes have the same view with respect to  $r$ ’s lock state.

### F. Correctness

By now, we have elaborated the details of MUREX. Below, we show the correctness of MUREX by the following theorems.

*Theorem 1 (Safety Property):* MUREX ensures the one-copy equivalence consistency criteria.

*Proof:* In MUREX, a read operation read-locks and accesses a read quorum of replicas, and returns the replica owning the largest version number. A write operation write-locks and updates a write quorum of replicas with the version number which is one more than the largest version number encountered. As shown in [13], multi-column quorums satisfy the intersection property: any pair of a read quorum and a write quorum, and any two write quorums have a non-empty intersection. By the read-write and write-write lock exclusions and the intersection property, a read operation always returns the replica with the largest version number, which is the most up-to-date. Furthermore, by the replica pointer, on-demand replica regeneration, and leased lock techniques, a replica can neither be locked by more than two write lock requests, nor be locked by a read lock request and a write lock request simultaneously. Thus, neither two write operations nor a write and a read operation can proceed at the same time. We can therefore conclude that one-copy equivalence is guaranteed. ■

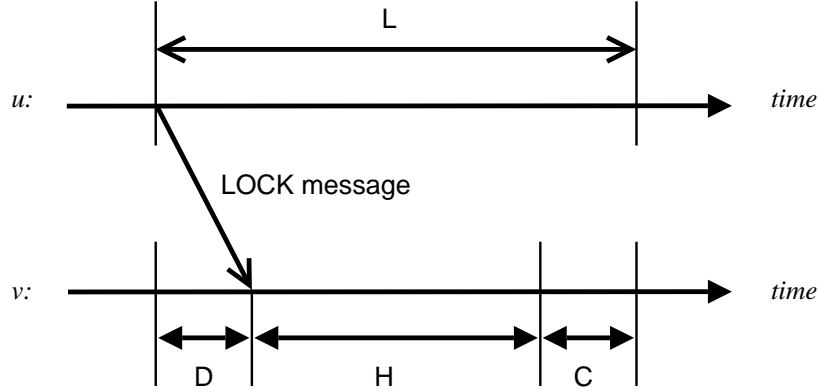
*Theorem 2 (Liveness Property):* There is neither deadlock nor starvation in MUREX.

*Proof:* MUREX demands a node to relinquish all locks when any lock is going to expire. This eliminates the condition of hold and wait, which is necessary for a deadlock to occur. Thus, there will be no deadlock. Furthermore, there will be no starvation since the random backoff mechanism demands a node to wait for a random delay before reissuing lock requests. Thus, a node can eventually acquire enough locks to perform the desired operation, which means that no starvation will occur. ■

## V. ANALYSIS AND SIMULATION

In this section, we analyze and simulate MUREX for evaluating its performance. We first analyze the communication cost of MUREX for the case of no contention. In MUREX, a node sends request messages for locking replicas in a last-column-to-first-column order. For an  $MC(m, s)$  multi-column structure, an operation needs  $3s$  messages in the best case: one LOCK, one OK and one UNLOCK messages sent to/by each member of the last column  $C_m$ .

When failures occur, the communication cost increases gradually. In the worst case, a node sends LOCK messages to all  $n$  replicas. Thus, the communication cost will be  $O(n)$  in the worst case. Fortunately, since MUREX demands a node to send LOCK messages to members of the last column first, the best case communication cost occurs much more frequently than the worst case.

Fig. 4. The relationship among  $H$ ,  $L$ ,  $C$  and  $D$ .

When there are contending nodes, the communication cost also increases. This is because when a node receives WAIT message, it will release all the locks obtained, and start over again to send LOCK messages after a random backoff time. Fortunately, the random backoff mechanism can scatter the re-sending of LOCK messages and thus some node will complete its operation successfully. Consequently, the communication cost will not go too high.

We conduct a simulation for MUREX with regard to success rates of operations for the purpose of evaluating the influence of different multi-column quorums. An operation is considered to be successful if it can finish before any leased lock expires. The simulation assumes that the underlying DHT is Tornado [12] and that there are totally 2000 nodes in the system. There are three experiments in our simulation. For each experiment, we perform the simulation by issuing 5000 operations, two thirds for reading and one third for writing. Each operation request is assumed to be destined for a random file (data object); thus, when the number of files increases, the degree of contention decreases.

In the first experiment, we assume there are 100 nodes that may join or leave the system randomly during the experiment. We also assume there are 2000 files published in the system and each node can accommodate at most 100 file replicas. We try to derive the success rate of multi-column structures  $MC(m, s)$  for  $m$  and  $s = 2..6$ . It is noted that when we simulate the case for multi-column structure  $MC(m, s)$ , the leased period is assumed to be  $m \times (\text{turn-around time})$ . Fig. 5 demonstrates the simulation results of the first experiment. By Fig. 5, we can observe that smaller  $s$  values render higher success rates. This is because only  $s$  replicas in the last column can form a quorum for an operation to proceed. By Fig. 5, we can also observe that more columns render lower success rates. This is because more columns imply more replicas, which in turn lead to higher probability that a node cannot accommodate all replicas. In the simulation, when a node cannot accommodate the replicas, it just drop the replica that is least recently used.

According to the above-mentioned observations, smaller  $s$  and smaller  $m$  are preferable for choosing multi-column structure  $MC(m, s)$ . However, when  $m = 2$ , we have the undesirable case that three out of four replicas are needed to

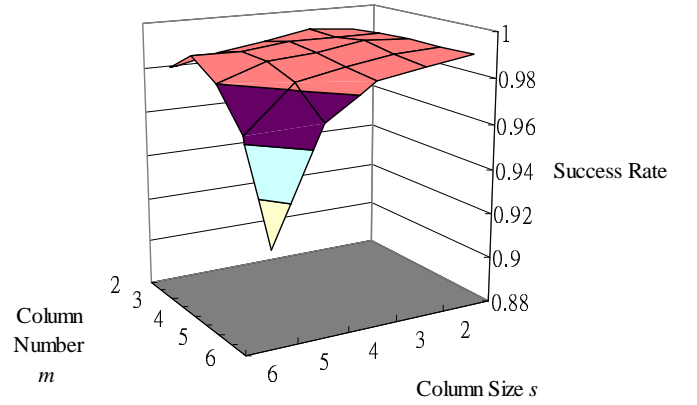


Fig. 5. Simulation results for the 1st experiment.

form a quorum. Therefore, we suggest to adopt multi-column structure  $MC(3, 2)$  for the construction of quorums. Actually, we adopt  $MC(3, 2)$  when we conduct the second and the third experiments.

The second experiment has the same parameter setting as the first experiment except that only multi-column structure  $MC(3, 2)$  is adopted and the number of files ranges from 50 to 3000. In Fig. 6, we can see that the success rate increases as the number of files increases. This is because the degree of contention decreases when there are more files. From this experiment, we can see that higher degrees of contention cause worse performances.

The third experiment has the same parameter setting as the first experiment except that only multi-column structure  $MC(3, 2)$  is adopted and the number of leaving nodes ranges from 200 to 1000. In Fig. 7, we can see that the success rate decreases as the number of leaving nodes increases. This is because more leaving nodes can cause more unsuccessful lock requests. From this experiment, we can see that higher node leaving rates cause worse performances.

## VI. RELATED WORK

As far as we know, there are four existent mutable P2P storage systems: Ivy [16], Eliot [21], Oasis [19], and Om [24]. The key concepts of these systems are P2P logs, replicated metadata service, dynamic quorum membership, and



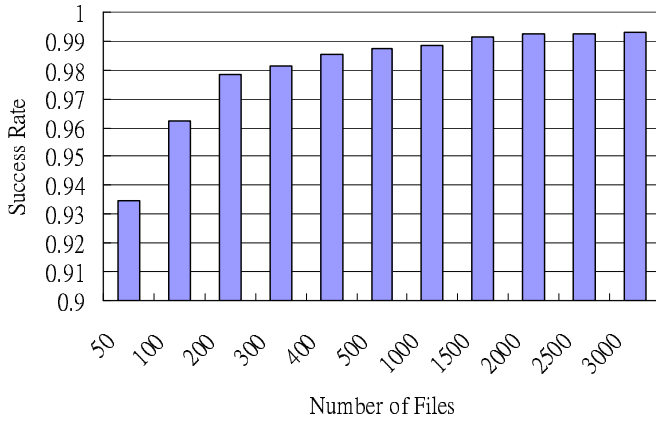


Fig. 6. Simulation results for the 2nd experiment.

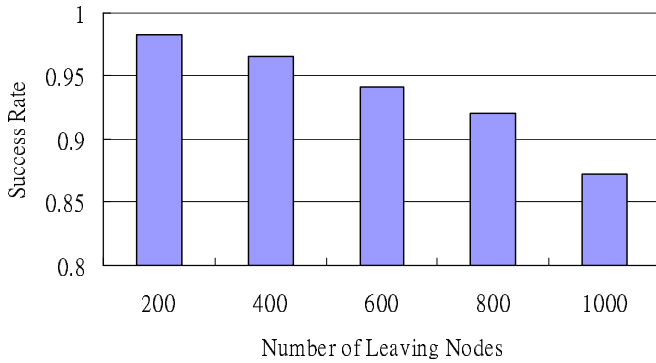


Fig. 7. Simulation results for the 3rd experiment.

replica membership reconfiguration, respectively. In addition to these four systems, we also review a scheme called “Sigma Protocol” [15], which intelligently collect replica states to achieve mutual exclusion among replicas. The scheme can be used as the basis of replica locking control and is thus worth investigating.

Below, we first introduce Ivy [16], which stores a set of logs with the aid of DHTs. Ivy keeps, for each participant, a log storing all its updates, and maintains data consistency optimistically by performing conflict resolutions among all logs. The logs should be kept indefinitely and a participant must scan all the logs associated with a file for the up-to-date data. Thus, Ivy is only suitable for small groups of participants.

Eliot [21] relies a reliable, fault-tolerant, and read-only P2P storage substrate — Charles to store data blocks, and uses an auxiliary metadata service (MS) for storing mutable metadata. It supports NFS-like write-through consistency semantics; however, the traffic between MS and the client is high for such semantics. It also supports AFS-like open-close consistency semantics; however, this semantics may cause the problem of lost updates. The MS service is provided by a conventional replicated database, which may be unfit for dynamic P2P environments.

Oasis [19] is based on Gifford’s weighted voting quorum concept and allows dynamic quorum membership. It spreads versioned metadata along with data replicas over the P2P network. To complete an operation on a data object, a client

must first find a metadata related to the object to figure out the total number of votes, required votes for a read and a write operation, replica list, and so on. It then forms a quorum according to the acquired metadata. Data consistency may be violated if a node happens to use a stale metadata.

Om [24] is based on the concepts of automatic replica regeneration and replica membership reconfiguration. The data consistency is maintained by two quorum systems: a read-one-write-all quorum system for accessing replicas, and a witness-based quorum system for reconfiguration. Om forwards all writes to the primary replica to serialize them, and uses a two-phase procedure to propagate the writes to all secondary replicas. In this manner, Om allows replica regeneration from a single replica. However, the primary replica may become a bottleneck and the overhead incurred by the two-phase procedure may be too high. Furthermore, the reconfiguration by witness-based quorum system has the probability of violating data consistency.

The Sigma protocol [15] uses Byzantine agreement algorithm [3], [14] to achieve mutual exclusion among replicas for P2P storage systems. The basic concept of the protocol is described as follows. A node  $u$  wishing to be the winner of the mutual exclusion should send a timestamped request to all  $n$  replicas ( $n = 3k + 1$ , where  $k$  stands for the maximum number of nodes that may leave the system) and waits for replies. On receiving a request from  $u$ , a node  $v$  should put  $u$ ’s request in a local queue in accordance with the timestamp order, takes the node as the winner whose request is in the front of the queue, and replies the winner ID to  $u$ . When the number of replies received by  $u$  exceeds  $2k + 1$ ,  $u$  acts according to the following conditions: (1) if more than  $2k + 1$  replies take  $u$  as the winner, then  $u$  is the winner. (2) if more than  $2k + 1$  replies take  $w$  ( $w \neq u$ ) as the winner, then  $w$  is the winner and  $u$  just keeps waiting. (3) if no node is regarded as the winner by more than  $2k + 1$  replies, then  $u$  sends YIELD message to cancel its request temporarily and then resends its request again. In this manner, one node can eventually be elected as the winner even when communication delay variance is large. However, Sigma protocol requires a node to send timestamped requests to all replicas and to receive advantaged replies from a large portion ( $\geq 2/3$ ) of replicas to be the winner of the mutual exclusion. This may incur large overhead; the overhead will even be larger under an environment of high contention.

## VII. CONCLUSION

In this paper, we have identified three problems for synchronous replication in DHT-based mutable P2P storage systems. The problems are replica migration, replica acquisition and state synchronization. We have proposed MUREX, a mutable replica control scheme, to solve these problems by the concepts of multi-column read/write quorums, replica pointers, on-demand replica regeneration and leased locks. We have proved that MUREX guarantees one-copy equivalence and causes no deadlock. Furthermore, we have analyzed and simulated MUREX. As we have shown, MUREX has constant communication cost in the best case and has good operation success rate.



## REFERENCES


- [1] D. Agrawal and A. E. Abbadi. The tree quorum protocol: An efficient approach for managing replicated data. In *16th International Conference on Very Large Data Bases, Brisbane, Queensland, Australia*, pages 243–254, 1990.
- [2] R. Bhagwan, D. Moore, S. Savage, and G. Voelker. Replication strategies for highly available peer-to-peer storage. In *Proc. of International Workshop on Future Directions in Distributed Computing*, pages 153–158, 2003.
- [3] M. Castro and B. Liskov. Practical Byzantine fault tolerance. In *Proc. of the Third Symposium on Operating Systems Design and Implementation*, 1999.
- [4] S. Y. Cheung, M. H. Ammar, and M. Ahamad. The grid protocol: A high performance scheme for maintaining replicated data. *IEEE Transactions on Knowledge and Data Engineering*, 4(6):582–592, 1992.
- [5] G. Chockler, D. Malkhi, and M. K. Reiter. Backoff protocols for distributed mutual exclusion and ordering. In *Proc. of the 21st International Conference on Distributed Computing Systems*, pages 11–20, 2001.
- [6] E. Cohen and S. Shenker. The gnutella 0.6 protocol draft, [http://rfc-gnutella.sourceforge.net/src/rfc-0\\_6-draft.html](http://rfc-gnutella.sourceforge.net/src/rfc-0_6-draft.html), June 2002.
- [7] E. Cohen and S. Shenker. Replication strategies in unstructured peer-to-peer networks. In *Proc. of SIGCOMM*, 2002.
- [8] F. Dabek, M. Kaashoek, D. Karger, R. Morris, and I. Stoica. Wide-area cooperative storage with CFS. In *Proc. of the 18th ACM Symposium on Operating Systems Principles (SOSP '01)*, 2001.
- [9] P. Druschel and A. Rowstron. Past: A large-scale, persistent peer-to-peer storage utility. In *Proc. of HotOS VIII*, pages 75–80, 2001.
- [10] D. K. Gifford. Weighted voting for replicated data. In *Proc. of 7th ACM symposium on Operating systems principles*, pages 150–162, 1979.
- [11] V. Gopalakrishnan, B. Silaghi, B. Bhattacharjee, and P. Keleher. Adaptive replication in peer-to-peer systems. In *Proc. of International Conference on Distributed Computing Systems*, pages 360–369, 2004.
- [12] H. C. Hsiao and C. T. King. Tornado: a capability-aware peer-to-peer storage overlay. *Journal of Parallel and Distributed Computing*, 64(6):747–758, 2004.
- [13] J. R. Jiang. The column protocol: A high availability and low message cost solution for managing replicated data. *International Journal of Information Systems*, 20(8):687–696, 1995.
- [14] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.
- [15] S. Lin, M. Lian, and Z. Zhang. A practical distributed mutual exclusion protocol in dynamic peer-to-peer systems. In *3rd International Workshop on Peer-to-Peer Systems (IPTPS'04)*, 2004.
- [16] A. Muthitacharoen, R. Morris, T. Gil, and B. Chen. Ivy: A read/write peer-to-peer file system. In *Proc. of 5th Symposium on Operating Systems Design and Implementation*, 2002.
- [17] Napster. Napster website, <http://www.napster.com>.
- [18] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Schenker. A scalable content-addressable network. In *Proc. of ACM SIGCOMM*, pages 161–172. ACM Press, 2001.
- [19] M. Rodrig and A. Lamarca. Decentralized weighted voting for P2P data management. In *Proc. of the 3rd ACM international workshop on Data engineering for wireless and mobile access*, pages 85–92, 2003.
- [20] A. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. In *Middleware*, pages 329–350, 2001.
- [21] C. Stein, M. Tucker, and M. Seltzer. Building a reliable mutable file system on peer-to-peer storage. In *Proc. of the 21st IEEE Symposium on Reliable Distributed Systems (SRDS'02)*, page 324, 2002.
- [22] I. Stoica, R. Morris, D. Karger, F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proc. of ACM SIGCOMM*, pages 149–160, 2001.
- [23] Robert H. Thomas. A majority consensus approach to concurrency control for multiple copy databases. *ACM Transactions on Database Systems (TODS)*, 4(2):180–209, 1979.
- [24] H. Yu and A. Vahdat. Consistent and automatic replica regeneration. *Trans. Storage*, 1(1):3–37, 2005.
- [25] B. Y. Zhao, J. Kubiatowicz, and A. Joseph. Tapestry: a fault-tolerant wide-area application infrastructure. *Computer Communication Review*, 32(1):81, 2002.




**Jehn-Ruey Jiang** received his Ph. D. degree in Computer Science in 1995 from National Tsing-Hua University, Taiwan, R.O.C. He joined Chung-Yuan Christian University as an Associate Professor in 1995. He joined Hsuan-Chuang University in 1998 and became a full Professor in 2004. He is currently with the Department of Computer Science and Information Engineering, National Central University. His research interests include distributed computing, mobile computing, peer-to-peer computing and pervasive computing.



**Chung-Ta King** received the B.S. degree in electrical engineering from National Taiwan University, Taiwan, R.O.C., in 1980, and the M.S. and Ph.D. degrees in computer science from Michigan State University, East Lansing, Michigan, in 1985 and 1988, respectively. From 1988 to 1990, he was an assistant professor of computer and information science at New Jersey Institute of Technology, New Jersey. In 1990 he joined the faculty of the Department of Computer Science, National Tsing Hua University, Taiwan, R.O.C., where he is currently a full professor. His current research interests include distributed processing and networked embedded systems. He is a member of the IEEE.



**Chi-Shiang Liao** received the B.S. and M.S. degrees in 2003 and 2005 from National Tsing Hua University in Hsinchu, Taiwan, R.O.C. His main research is on Peer-to-Peer Networks and Pervasive Computing. He is working in Infotrend Corporation in Taiwan, R.O.C.



**Ching-Hao Liu** received the B.S and M.S. degrees in computer science from National Chiao Tung University and National Central University, Taiwan, R.O.C., in 2001 and 2003, respectively. Now, he is a Ph.D. candidate in National Tsing Hua University, Taiwan, R.O.C. His research interests include distributed systems and P2P computing.