Replication Requirements in Mobile Environments*

David Ratner Peter Reiher Gerald J. Popek†
Department of Computer Science
University of California, Los Angeles

Abstract

Replication is required in mobile environments because nomadic users require local copies of important data. However, today's replication systems are not "mobile-ready." Instead of enabling mobile computing and improving the mobile user's environment, the replication system actually hinders mobility and complicates mobile operation. The replication services are designed for stationary environments, and as such do not and cannot provide mobile users with the capabilities they require. Replication in mobile environments requires fundamentally different solutions than those previously proposed, because nomadicity presents a fundamentally new and different computing paradigm. Here we outline the key requirements that mobility places on the replication service, and briefly describe Roam, a system designed to meet those requirements.

1 Introduction

Mobile computing is rapidly becoming standard in all types of environments: academic, commercial, and private. Widespread mobility impacts multiple arenas, but one of particular importance is data replication. Replication is especially important in mobile environments, since disconnected or poorly connected machines must rely primarily on local resources. The monetary costs of communication when mobile, combined with the lower bandwidth, higher latency, and reduced availability, effectively require that important data be stored locally on the mobile machine. In the case of shared data, between multiple mobile users or between mobile and stationary machines, replication is often the best and sometimes only viable approach.

Many replication solutions [3, 14] assume a static infrastructure; that is, the connections themselves may be transient but the connection *location* and the set of possible synchronization partners always remain the same. However, mobile users are by definition not static, and a replication service that forces them to adjust to a static infrastructure hinders mobility rather than enables it. Extraordinary actions, such as long distance telephone calls over low-bandwidth links, are necessary for users to conform to the underlying static model, costing additional time and money while providing a degraded service. Additionally, mobile users have difficulty inter-operating with other mobile users, because communication patterns and topologies are typically predefined according to the underlying infrastructure. Often, direct synchronization between mobile users is simply not permitted.

Other systems [1, 12, 16] have simply traded the above communication problem for another one: scaling. They provide the ability for *any-to-any* synchronization, but their model suffers from inherent scaling problems, limiting its usability in real environments. Good scaling behavior is very important in the mobile scenario. Mobile users clearly require local replicas on their mobile machines. Yet, replicas must also be stored in the office environment for reliability, intra-office use by non-mobile personnel, and system-administration activities like back-ups. Additionally, typical methods for reducing replication factors, such as local area network sharing techniques, are simply not feasible in the mobile context. Mobile users require local replicas of critical information, and in most cases desire local access to non-critical objects as well for cost and performance reasons. The inability to scale well is as large an obstacle to the mobile user as the insufficient communication model discussed above.

The main problem is that mobile users are replicating data using systems that were not designed for mobility. As such, instead of the replication system improving the state of mobile computing, it actually hinders mobility, as users find themselves forced to adjust their physical motion and computing needs to better match what the system expects. This paper outlines the requirements of a replication service designed for

*This work was sponsored by the Advanced Research Projects Agency under contract DABT63-94-C-0080. The authors can be reached at the Department of Computer Science, UCLA, Los Angeles, CA 90095, or by email to {ratner,popek,reiher}@cs.ucla.edu.

†Gerald Popek is also affiliated with Platinum *technology*.

the mobile context. We conclude with a description of Roam, a replication solution redesigned especially for mobile computing. Built using the Ward architecture [10], it enables rather than hinders mobility, and provides a replication environment truly suited to mobile environments.

2 Replication Requirements

Mobile users have special requirements above and beyond those of simple replication required by anyone wishing to share data. Here we discuss some of the requirements that are particular to mobile use: any-to-any communication, larger replication factors, and detailed controls over replication behavior. We omit discussion of well-understood ideas, such as the case for optimistic replication, discussed in [1, 2, 4, 15].

2.1 Any-to-any communication

By definition, mobile users change their geographic location. As such, it cannot be predicted *a priori* what machines will be geographically co-located at any given time. Given that it is typically cheaper, faster, and more efficient to communicate with someone local rather than remote, mobile users want the ability to directly communicate and synchronize with whomever is "nearby," primarily for cost reasons. Consistency can of course be correctly maintained even if two machines cannot directly synchronize with each other, as demonstrated by systems based on the client-server model [3, 14]. The question is one of usability, expected functionality, and inherent cost. Users who are geographically co-located don't want updates to *eventually* propagate through a long-distance, sub-optimal path; the two machines are next to each other, and the synchronization should be instantaneous.

Since users expect that nearby machines can synchronize with each other quickly and efficiently, and it cannot be predicted which machines will be geographically co-located at any point in the future, a replication model capable of supporting *any-to-any communication* is required. That is, the model must allow any machine to communicate with any other machine; there can be no second-class clients in the system.

Any-to-any communication is also required in other mobile arenas, such as in *appliance* mobility [5], the motion from device to device or system to system. For instance, given a desktop, a laptop, and a palmtop, it is unlikely that one would want to impose a strict client-server relationship between the three; rather, one would want them all to be able to communicate with each other.

Providing any-to-any communication is equivalent to utilizing a *peer-to-peer* replication model [9, 12, 16]; if anyone can directly synchronize with anyone else, then everyone must by definition be equals, at least with respect to update-generation abilities. Some, however, have argued against peer models in mobile environments because of the relative insecurity regarding the physical devices themselves—laptops are often stolen at airports and hotels. The argument is that since mobile computers are physically less secure, they should be "second-class" citizens with respect to the highly secure servers located behind locked doors [13]. The class-based distinction supposedly provides improved security by limiting the potential security breach to only a second-class object.

The argument is based on the assumption that security features must be encapsulated within the peer model, and therefore the unauthorized access of any peer thwarts all security barriers and mechanisms. Fortunately, the assumption is not true. Systems such as Truffles [11] have demonstrated that security policies can be modularized and logically situated around a peer replication framework while still remaining independent of the replication system. With such an architecture, the problems caused by the unauthorized access of a peer replica are no different from the unauthorized access of a client in a client-server model. Thus, the question of update-exchange topologies (any-to-any as compared to a more stylized, rigid structure as in client-server models) can be dealt with independently of the security issue and the question of how to enforce proper security controls.

2.2 Larger replication factors

Most replication systems only provide for a handful of replicas of any given object. Additionally, peer algorithms have never traditionally scaled well. Finally, some have argued that peer solutions simply by their nature cannot scale well [13].

However, while mobile environments seem to require a peer-based solution (described above), they also seem to negate the assumption that a handful of replicas is enough. While we do not claim a need for thousands of writable copies, it does seem likely that the environments common today and envisioned for the near future will require larger replication factors than current systems allow. First and foremost, each mobile user requires a local replica on their laptop, doubling replication factors when data is stored both on the users desktop and laptop. Additionally, replication factors can often be minimized in office environments due to LAN-style sharing and remote-access capabilities; one

machine stores a replica that every machine in the office can access quickly, cheaply and easily. However, network-based file sharing cannot be utilized in mobile environments due to the frequency of network partitions and the variable ranges of available bandwidth and transfer latency.

Second, consider the case of appliance mobility. The above discussion assumes that each user has one static machine and one mobile machine. The future could potentially see the use of many more "smart" devices capable of storing replicated data. Palmtop computers are becoming more common, and there is even a wristwatch that can download calendar data. Various researchers [17] have built systems that allow laptop and palmtop machines to share data dynamically and opportunistically. It is not difficult to imagine other devices in the near future having the capability to store and conceivably update replicated data; such devices potentially increase replication factors dramatically.

Finally, some have argued the need for larger replication factors independent of the mobile scenario, such as in the case of air traffic control [8]. Other scenarios possibly requiring larger replication factors include stock exchanges, routing, airline reservation systems, and military command and control.

Read-only strategies and other class-based techniques cannot adequately solve the scaling problem, at least in the mobile scenario. Class-based solutions are not applicable to mobility, for the reasons described above (Section 2.1). Read-only strategies are not viable solutions because they force users to pre-select the writable replicas beforehand and limit the actual number of writable copies. In general one cannot predict which replicas require write-access and which ones do not; we must provide the *ability* for all replicas to generate updates, even though some may never do so.

2.3 Detailed replication controls

A replication service by definition provides users with some degree of replication *control*—a method of indicating what objects they want replicated. Many systems provide replication on a large-granularity basis, meaning that users requiring one portion of the container must locally replicate the entire container. Such systems are perhaps adequate in stationary environments, when users have access to large disk pools and network resources, but replication control becomes vastly more important to mobile users. Nomadic users do not in general have access to off-machine resources, and therefore objects that are not locally stored are effectively inaccessible. Thus, everything the user requires must be replicated locally, which becomes problematic when the container is large.

Replicating a large-granularity container means that some of the replicated objects will be deemed unimportant to the particular user. Unimportant data occupies otherwise usable disk space, which cannot be used for more critical objects. In the mobile context, where network disconnections are commonplace, important data that cannot be stored locally causes problems ranging from minor inconveniences to complete stoppages of work and productivity, as described by Kuenning [7]. Kuenning's studies of user behavior indicate that the set of required data can in fact be completely stored locally, but only if the underlying replication service provides the appropriate flexibility to individually select objects for replication. Users and automated tools therefore require fairly detailed controls over what objects are replicated, because without them mobile users cannot adequately function.

3 Roam

Roam is a system designed to meet the above set of requirements. It is based on the *Ward* model [10] and is currently being developed at the University of California at Los Angeles.

3.1 Ward model

The Ward model combines classical elements of both the traditional peer-to-peer and client-server models, yielding a solution that scales well and provides replication flexibility, allowing dynamic reconfiguration of the synchronization topology. The model's main grouping mechanism is the *ward*, or Wide Area Replication Domain. A ward is a collection of "nearby" machines, geographically co-located in some area and possibly only loosely connected. While all members of the ward are equal peers, the ward has a designated *ward master*, similar to a server in a client-server model but with several important differences:

- Since all ward members are peers, any two ward members can directly synchronize with one another. Typical client-server solutions do not allow client-client synchronization. Whether by design or by accident, mobile users will encounter other mobile users; in such cases, direct access to the other ward member may be easier, cheaper and more efficient than access to the ward master.

- Since all ward members are peers, any ward member can serve as the ward master. Automatic reelection and *ward-master reconfiguration* can occur should the ward master fail or become un-

available, and algorithms exist to resolve multiple-master scenarios. Correctness is not affected by a transient ward master failure, but the system maintains better consistency if the ward master is typically available and accessible.

- The ward master is not required to store actual data for all intra-ward objects, though it must be able to identify (i.e. name) the complete set. Most client-server strategies force the server to store a superset of each client's data.

The ward master is the ward's only link with other wards; that is, only the ward master is aware of other replicas outside the ward. This is one manner in which the ward model achieves good scaling—by limiting the amount of knowledge stored at individual replicas. Traditional peer models force every replica to learn about other replica's existence; in the ward model, replicas are only knowledgeable about the other replicas within their ward.

All ward masters belong to a higher-level ward, forming a two-level hierarchical model.[1] Ward masters act on their ward's behalf by bringing new updates into the ward, exporting others out of the ward, and gossiping about all known updates. Consistency is maintained across all replicas by having ward masters communicate with each other and allowing information to propagate independently within each ward. Figure 1 illustrates the basic architecture, as well as advanced features discussed in later sections.

Wards are dynamically formed at replica-creation time, and are dynamically maintained as suitable ward-member candidates change. Ward destruction occurs automatically when the last replica in a given ward is destroyed.

3.2 Flexibility in the model

Replication flexibility is an important feature of the ward model. The set of data stored within each ward, called the *ward set*, is dynamically adjustable, as is the set of ward members themselves. As ward members change their data demands and alter what replicated data they store locally, the ward set changes. Similarly, as mobile machines join or leave the ward, the set of ward participants changes. Both the ward set and ward membership are locally recorded and are replicated in an optimistic fashion.

Additionally, each ward member, including the ward master, can locally store a different subset of the ward

[1] The rationale behind the two-level hierarchy and its impact on scaling is discussed in Section 3.4.

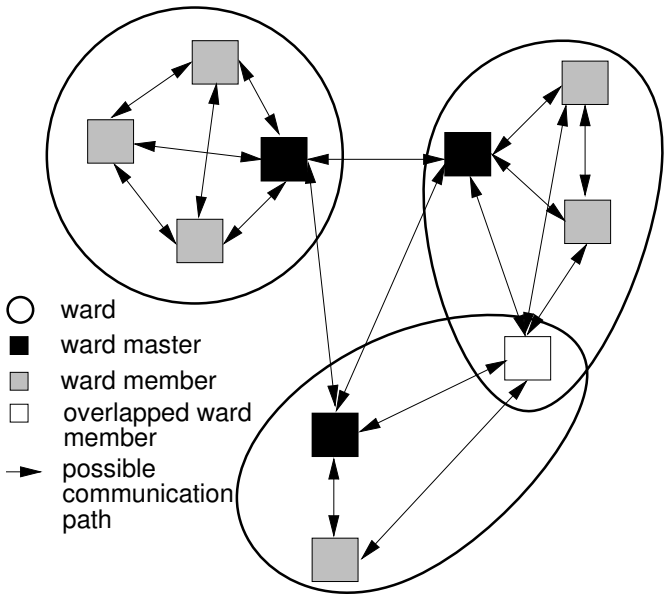

Figure 1: The basic ward architecture. Overlapped members are a mobility feature (Section 3.3).

set. Such replication flexibility, called *selective replication* [9] provides improved efficiency and resource utilization: ward members locally store only those objects that they actively require. Replication decisions can be made manually or with automated tools [4, 6].

Since the ward set varies dynamically, different wards might store different sets: not all ward sets will be equivalent. The model in essence provides selective replication between wards themselves. The reconciliation topologies and algorithms [9] apply equally well within a single ward and between ward masters. Briefly, the algorithms provide that machines communicate with multiple partners to ensure that each data object is synchronized directly with another replica. Additionally, the data synchronization algorithms support the *reconciliation* of non-local data via a third-party data-storage site, allowing the ward master to reconcile data which it does not locally store but is stored somewhere within the ward.

3.3 Support for mobility

The model supports two types of mobility. *Intra-ward* mobility occurs when machines within the same ward become mobile within a limited geographic area; the machines encountered are all ward members. Since ward members are peers, direct communication is possible with any encountered machine. Intra-ward mobility might occur within a building, when traveling to a co-worker's house, or at a local coffee shop.

Perhaps more interesting, *inter-ward* mobility occurs when users travel (with their data) to another geographic region, encountering machines from another ward. Examples include businessmen traveling to remote offices and distant collaborators meeting at a common conference.

Inter-ward mobility raises two main issues. First, recall that due to the model's replication flexibility, two wards might not have identical ward sets. Thus, the mobile machine may store data objects not kept in the new ward, and vice-versa. Second, consider the typical patterns of mobility. Often users travel away from their "home location" for only a short time. The system would perform poorly if such transient mobile actions required global changes in data structures across multiple wards. On the other hand, mobile users occasionally spend long periods of time at other locations, either permanently or semi-permanently changing their definition of "home." In these scenarios, users should be provided with the same quality of service (in terms of local performance and time to synchronize data) as they experienced in their previous "home".

Our mobility solution resolves both issues by defining two styles of inter-ward mobility—short-term (transient) and long-term (semi-permanent)—and providing the ability to transparently and automatically upgrade from the former to the latter. The two operations are called *ward overlapping* and *ward changing* respectively. Collectively, the two are called *ward motion* and enable peer-to-peer communication between any two replicas in the ward model, regardless of their ward membership.

3.3.1 Ward changing

Ward changing involves a long-term, perhaps permanent, change in ward membership. The moving replica physically changes its notion of its "home" ward, forgetting all information from the previous ward; similarly, the other participants in the old and new wards alter their notion of current membership. Since ward membership information is maintained optimistically, the problem of tracking membership in often-disconnected environments is straightforward.

The addition of a new ward member may change the ward set. Since the ward master is responsible for the inter-ward synchronization of all data in the ward set, the ward set must expand to properly encompass the replicated data stored at the moving replica. Similarly, the ward set at the old ward may shrink in size, as the ward set is dynamically and optimistically recalculated when ward membership changes. The ward-set changes propagate to other ward masters in an optimistic, "need-to-know" fashion so that only the ward masters that care about the changes learn of them. Since both ward sets potentially change, and these changes are eventually propagated to other ward masters, ward changing is a heavyweight operation. However, users benefit because all local data can be synchronized completely within the local ward, giving users the best possible quality of service and reconciliation performance.

3.3.2 Ward overlapping

In contrast, ward overlapping is intended as a very lightweight mechanism, and causes no global changes within the system. Only the new ward is effected by the operation. The localization of changes makes it a lightweight operation both to perform and to undo.

Ward overlapping allows simultaneous multi-ward membership, enabling direct communication with the members of each ward. To make the mechanism lightweight, we avoid changing the ward sets by making the new replica an "overlapped" member instead of a full-fledged participant. Ward members (except for the ward master) cannot distinguish between real and overlapped members; the only difference is in the management of the ward set. Instead of merging the existing ward set with the data stored on the mobile machine, the ward set remains unaltered. Data shared between the mobile machine and ward set can be reconciled locally with members of the new ward. However, data outside the new ward set cannot be reconciled locally, and must either temporarily remain unsynchronized or else be reconciled with the original home ward.

3.3.3 Ward motion summary

When a replica enters another ward, there are only two possibilities: the ward set can change or remain the same. The former creates an performance-improving but heavyweight solution; the latter causes a moderate performance hit (when synchronizing data not stored in the new ward) but provides a very lightweight solution for transient mobile situations. Since both are operationally equivalent, the system can transparently upgrade from overlapping to changing if the motion seems more permanent than first expected.

Additionally, since ward formation is itself dynamic, users can easily form *mobile workgroups* by identifying a set of mobile replicas as a new (possibly temporary) ward. Mobile workgroups can be formed without leaving the old ward by using ward overlapping. Ward motion and dynamic ward formation and destruction allow easy and straightforward communication between

any set of replicas in the entire system.

3.4 Scalability

The scalability of the ward model is directly related to the degree of replication flexibility. Ward sets can dynamically change in unpredictable ways; therefore the only method for a ward master to identify its ward set is to list each entry individually. The fully hierarchical generalization of the ward model to more than two levels faces scaling problems due to the physical problems of storing and indexing these lists of entries.

Nevertheless, the proposed model scales well within its intended environment, and allows several hundred read-write replicas of any given object, meeting the demands of everyone from a single developer or a medium-sized committee to a large, international company. The model could be adapted to scale better by restricting the degree of replication freedom. For instance, if ward sets changed only in very regular fashions, they could be named as a unit instead of naming all members, dramatically improving scalability. However, we believe that replication flexibility is an important design consideration in the targeted mobile environment, and one that users absolutely require.

4 Conclusion

Replication is required for mobile computing, but today's replication services do not provide the key features required by mobile users. Nomadicity requires a replication solution that provides any-to-any communication in a scalable fashion with sufficiently detailed control over the replication decisions. ROAM is a system being designed and implemented to meet these goals, paving the way not just to improved mobile computing but to new and better avenues of mobile research.

References

[1] R. G. Guy, J. S. Heidemann, W. Mak, T. W. Page, Jr., G. J. Popek, and D. Rothmeier. Implementation of the Ficus replicated file system. In *USENIX Conference Proceedings*, pages 63–71. University of California, Los Angeles, USENIX, June 1990.

[2] J. S. Heidemann, T. W. Page, Jr., R. G. Guy, and G. J. Popek. Primarily disconnected operation: Experiences with Ficus. In *Proceedings of the Second Workshop on Management of Replicated Data*, pages 2–5. University of California, Los Angeles, IEEE, Nov. 1992.

[3] P. Honeyman, L. Huston, J. Rees, and D. Bachmann. The Little Work project. In *Proceedings of the Third Workshop on Workstation Operating Systems*, pages 11–14. IEEE, Apr. 1992.

[4] J. J. Kistler and M. Satyanarayanan. Disconnected operation in the Coda file system. *ACM Transactions on Computer Systems*, 10(1):3–25, 1992.

[5] L. Kleinrock. Nomadicity. Presentation at the GloMo PI Meeting (February 4) at the University of California at Los Angeles, 1997.

[6] G. H. Kuenning. Design of the SEER predictive caching system. In *Proceedings of the Workshop on Mobile Computing Systems and Applications*, Santa Cruz, CA, Dec. 1994.

[7] G. H. Kuenning. *Seer: Predictive File Hoarding for Disconnected Mobile Operation.* PhD thesis, University of California, Los Angeles, Los Angeles, CA, May 1997. Also available as UCLA CSD Technical Report UCLA-CSD-970015.

[8] P. Quéinnec and G. Padiou. Flight plan management in distributed air traffic control system. In *Proceedings of the International Symposium on Autonomous Decentralized Systems*, Kawasaki, Japan, Mar. 1993.

[9] D. Ratner, G. J. Popek, and P. Reiher. Peer replication with selective control. Technical Report CSD-960031, University of California, Los Angeles, July 1996.

[10] D. Ratner, G. J. Popek, and P. Reiher. The ward model: A scalable replication architecture for mobility. In *Workshop on Object Replication and Mobile Computing*, Oct. 1996.

[11] P. Reiher, T. Page, S. Crocker, J. Cook, and G. Popek. Truffles—a secure service for widespread file sharing. In *Proceedings of the The Privacy and Security Research Group Workshop on Network and Distributed System Security*, Feb. 1993.

[12] P. Reiher, J. Popek, M. Gunter, J. Salomone, and D. Ratner. Peer-to-peer reconciliation based replication for mobile computers. In *Proceedings of ECOOP'96 II Workshop on Mobility and Replication*, July 1996.

[13] M. Satyanarayanan. The influence of scale on distributed file system design. *IEEE Transactions on Software Engineering*, SE-18(1):1–8, Jan. 1992.

[14] M. Satyanarayanan, J. J. Kistler, P. Kumar, M. E. Okasaki, E. H. Siegel, and D. C. Steere. Coda: A highly available file system for a distributed workstation environment. *IEEE Transactions on Computers*, 39(4):447–459, Apr. 1990.

[15] M. Satyanarayanan, J. J. Kistler, L. B. Mummert, M. R. Ebling, P. Kumar, and Q. Lu. Experience with disconnected operation in a mobile computing environment. In *Proceedings of the USENIX Symposium on Mobile and Location-Independent Computing*, pages 11–28, Cambridge, MA, Aug. 1993. USENIX.

[16] D. B. Terry, M. M. Theimer, K. Petersen, A. J. Demers, M. J. Spreitzer, and C. H. Hauser. Managing update conflicts in Bayou, a weakly connected replicated storage system. In *Proceedings of the 15th Symposium on Operating Systems Principles*, pages 172–183, Copper Mountain Resort, Colorado, Dec. 1995. ACM.

[17] R. Want, B. N. Schilit, N. I. Adams, R. Gold, K. Petersen, D. Goldberg, J. R. Ellis, and M. Weiser. An overview of the ParcTab ubiquitous computing experiment. *IEEE Personal Communications Magazine*, 2(6):28–43, Dec. 1995.