

# Creating and maintaining replicas in unstructured peer-to-peer systems

*Elias Leontiadis\* Vassilios V. Dimakopoulos Evaggelia Pitoura*

Technical Report 2006-1

Department of Computer Science, University of Ioannina  
P.O. Box 1186, Ioannina, Greece GR-45110

Jan. 2006

## Abstract

Peer-to-peer systems have received much attention as a way of sharing data among a large and dynamic population of nodes. In such systems replication is an important issue as it improves search performance and data availability. It has been shown that optimal replication is attained when the number of replicas per item is proportional to the square root of their popularity. In this paper, we focus on updates in the case of optimal replication. In particular, we propose a new practical strategy for achieving square root replication called pull-then-push replication (PtP). With PtP, after a successful search, the requesting node enters a replicate-push phase where it transmits copies of the item to its neighbors. We show that updating the replicas can be significantly improved through an update-push phase where the node that created the copies propagates any updates it has received using similar parameters as in replicate-push. Our experimental results show that replicate-push coupled with an update-push strategy achieves good replica placement and consistency with small message overhead.

---

\*This work was done mainly while E. Leontiadis was with the University of Ioannina. He is currently with the Department of Computer Science, University College London, United Kingdom.



## 1. Introduction

The popularity of file sharing systems (such as Napster and Gnutella) has resulted in attracting much current research in peer-to-peer (p2p) systems. Peer-to-peer systems offer a means for sharing data among a large, diverse and dynamic population of users. An issue central in such systems is resource location, i.e. given a user query for data, to discover the peers with matching data items.

There are two basic approaches for building p2p systems for efficiently locating data. In structured p2p systems, data items are assigned to specific peers using some form of distributed hashing. Locating peers with matching data is then guaranteed to take place by visiting a bounded number of peers, normally logarithmic to the total number of peers in the system. In unstructured p2p systems, there is no assumption about the placement of data items. New nodes connect to some other nodes in the p2p system randomly. When compared with structured p2p systems, unstructured p2p systems usually provide no guarantees for search performance but do not suffer from the cost induced from maintaining the structure and from load balancing procedures necessary in structured p2p systems.

In this paper, we focus on the problem of replication in unstructured p2p systems. Replication improves the performance of search as well as data availability. Availability issues are especially critical in p2p systems, since peers leave the system very often, thus making their data unavailable.

Previous work on the topic [1, 8] has shown that optimal (with respect to search performance) replication is achieved when the number of copies per data item is proportional to the square root of their popularity. In this paper, we propose a new practical strategy for achieving square root replication called pull-then-push replication (PtP). With PtP replication, after a successful search for a data item, the node that posed the query enters a replicate-push phase during which it pushes copies of the item to its neighbors.

Then, we propose consistency maintenance protocols for copies created using the optimal replication strategy. We show that updating the copies can be significantly improved through an update-push phase where the node that created the copies propagates any updates it receives to its neighbors. Although, replica consistency

protocols have been previously proposed (e.g., in [2]), our main contribution is that we study the problem in conjunction with the strategy used to create the copies. Our experiments show that the best results are achieved when update-push uses similar parameters with replicate-push.

The remainder of this paper is structured as follows. In Section 2, we summarize previous work on optimal replication and we present our own PtP strategy for achieving it. We also evaluate its performance experimentally. In Section 3, we introduce our consistency maintenance protocols, along with experimental results. Finally, Section 4 concludes the paper.

## 2. Optimal replication

Suppose there are in total  $m$  different data items in the network, and that, collectively, the peers have capacity for storing  $R$  items<sup>1</sup>. Also, assume that the query rate for item  $i$  is  $q_i$ ,  $i = 1, \dots, m$ . Cohen and Shenker [1] developed a theory for optimally replicating the data items in unstructured peer-to-peer networks, given the restriction of  $R$ . In particular, they studied different replication strategies and showed that the expected search cost is minimized when the  $i$ th item has  $r_i$  replicas, where  $r_i$  is proportional to  $\sqrt{q_i}$ .

In their analysis, the authors assumed a theoretical random probes (RP) search method: the inquiring node repeatedly probes peers in random and asks for the item, until the item is found. As the authors argued, the RP method captures the essential behavior of the blind search strategies usually employed in p2p systems (such as flooding) because in unstructured networks the topology is unrelated to the location of data.

The problem with square-root (SR) replication is that it requires knowledge of the query rate for each item. To alleviate this, the following scheme was proposed: after each successful search, the item is copied to a number of nodes equal to the number of probes. It can be shown that, assuming an analogous rate of item removals, this scheme leads to SR replication.

---

<sup>1</sup>Data items can be actual copies of the data or just pointers to them.

However, even this scheme is not easily implementable. Keeping track of the number of queried nodes is simply impractical when the usual flooding-based search algorithms are used, due to the excessive number of messages required. But even if a practical way of counting the queried nodes existed, this number would not be equal to the number of random probes that would have been required. The reason is that the theoretic RP strategy stops immediately after locating the item. All practical strategies, however, unleash parallel search paths — if the item is found in one of the search paths, the rest might continue querying nodes until, for example a time-to-live (TTL) parameter was exhausted.

In conclusion, practical strategies for approximating the number of probes are required. In [8], the authors examined a number of such algorithms, namely *owner-replication*, *path-replication* and *random-replication*. In owner-replication, the inquiring node is the only one that makes a copy of the resource — leading clearly to suboptimal replication. In the other two strategies, the node that provides the resource creates a number of replicas, equal to the distance (in hops) between the inquiring and the offering node. The only difference between the last two strategies is where the replicas are placed. Path and random replication approach SR replication but not quite accurately. The reason is that if the distance between the inquiring and the offering node is  $t$  hops, the RP strategy may not have located the item within just  $t$  probes, unless a single path was used for the search. The authors used multiple random walkers, which naturally visit a multiple of  $t$  nodes. We next propose a simple but effective scheme.

## 2.1. Pull-Then-Push replication

The proposed scheme is based on the following idea: the creation of replicas is delegated to the *inquiring* node, not the providing node. The scheme consists of two phases. The *pull* phase refers to searching for a data item. After a successful search, the inquiring node enters a *push* phase, whereby it transmits the data item to other nodes in the network in order to force creation of replicas for that item. We call this the *Pull-then-Push* (PtP) replication. One can conceive variations of the PtP strategy by utilizing different algorithms for the pull and push phases. Path

replication as suggested in [8] could be considered as a type of PtP replication, where the pull phase uses multiple random walkers, while the push phase uses a single path.

In order to reach SR replication, we need to create a number of replicas equal to the number of probed nodes<sup>2</sup>. Consequently, one should utilize the *same* algorithm for the push and the pull phases, so that the push phase visits approximately the same nodes the pull phase visited. Thus, for example, if a random BFS [4] or teeming [3] search algorithm is used for the pull phase, the same algorithm should be used to broadcast the item during the push phase.

All practical search strategies produce multiple search routes, and utilize some form of TTL to limit the search space (and the resulting message overhead). If during the pull phase the item was found at distance  $t$  hops from the inquirer, then the push phase should also stop after  $t$  hops. This means that the TTL utilized for the push phases should not be set according to the TTL used during pull, but rather according to  $t$ .

However, because of the multiple search routes produced, the  $t$ th step may contact quite a large number of nodes. In [3], it was shown that for pure flooding, the number of messages grows exponentially with the TTL; most of those messages are sent in the last step of the search. For example, assume a random network with each peer connected to  $d$  other nodes, and a pure flooding strategy, where each peer propagates the query to all its neighbors. If a search returned an item at the 3rd step, approximately  $d + d^2 + d^3$  different peers would have been visited, although only one node at distance 3 had the item. This means that  $d + d^2 + 1$  probes could be enough and as a result, the best strategy for the push phase would be to use a TTL of 2, not 3. In general, the TTL used for the push phase should be equal to the hop distance at which the item was found minus one.

Recapping, our proposed PtP strategy adheres to the following rules:

- After a successful search, the requester pushes the item back to the network.
- The same algorithm is used for both pull (search) and push.
- The TTL for push is equal to  $t - 1$ , where  $t$  is the hop distance where the

---

<sup>2</sup>To be more precise, equal to the number of probes the RP search strategy would require.

resource was found.

- All peers receiving the push message create a replica of the item.

In the next section we provide simulation results which confirm that this simple PtP strategy does indeed lead to square-root replication.

## 2.2. Experimental results

The PtP strategy has been evaluated through extensive simulations. In our simulator, we construct a network of peers/nodes, where each peer is connected to  $d$  other peers in random, called its neighbors. Each peer offers a number of data items and also has a fixed number of slots for replicating other items. Initially, all replica slots are empty. Then, we continuously perform searches originating at random peers, for random items. After each search, a push phase occurs, where replication is forced according to the strategy used. If a peer has to replicate an item and has no available slot, a uniformly random slot is emptied so that room is created for the new replica. Results are collected after a sufficiently large number of searches; the single most important metric we extract is the number of replicas,  $r_i$ , for each item.

The simulator is capable of utilizing a number of different search (pull) strategies. In all these strategies, a peer that receives a query for a data item, first checks whether it knows about the item; if not, it propagates the query to its neighbors. The strategies differ in the set of neighbors where the queries are propagated, and include:

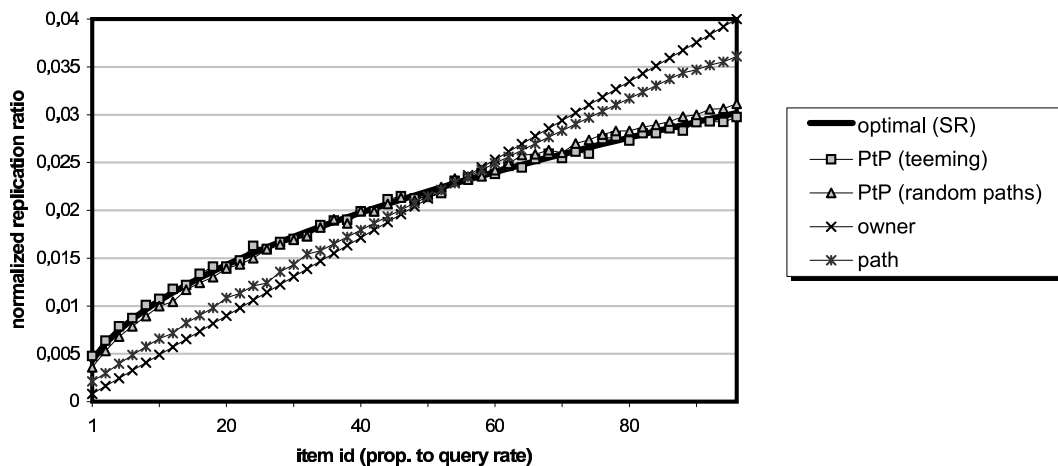
- *Pure flooding.* Peers propagate the query to all their neighbors. This is the strategy with the highest message overhead and the poorest scalability.
- *Random walkers or random paths* [8]. Each peer propagates the query to exactly one of its neighbors, in random. A single random path incurs quite few message transmissions but has the disadvantage of requiring quite a large number of hops before locating the data item. Multiple walkers searching in parallel is a variation to decrease the average number of hops: the inquiring node sends the query to a number of its neighbors, each one unleashing a random walker.

- *Random BFS or teeming* [4, 3]. Peers propagate the query to each of their neighbors with some fixed probability  $\phi$ . For  $\phi = 1$  we have pure flooding. Smaller values of  $\phi$  produce a smaller message count, at the expense of an increased number of steps to locate an item.
- *Teeming with decay* [7]. This works like teeming but the probability  $\phi$  decreases with the distance from the inquiring node, so as to avoid the high message count in later steps. If a node is in distance  $t$  from the inquiring peer, then the probability is given by:

$$\phi_t = \phi(1 - c)^t,$$

where  $\phi_0 = \phi$  and  $c$  is the *decay* parameter. For  $c = 0$  we have simple teeming, while if in addition  $\phi = 1$ , the strategy is pure flooding. The higher the value of  $c$ , the faster the probability decreases and the fewer the message transmissions become.

The same algorithms are used for the push phase. Of course, in this case the peers do not receive queries but just messages/items which they propagate immediately to some of their neighbors, according to the method used.



**Figure 1.** *Distribution of replication ratios under various replication strategies*

In Fig. 1, we present results for a network of 1000 peers where each one has storage space for 10 items, out of a total of  $R = 100$  different items. The replication strategies employed are owner, path, and PtP replication. For PtP, we experimented



with all the algorithms presented above and with different parameters. In Fig. 1, we show the results for two of them, one with 5 random walkers and  $TTL = 10$  and one with teeming,  $TTL = 5$  and a decay parameter of  $c = 0.4$ . The other algorithms exhibited the same behavior, and were omitted for clarity.

The plot shows the normalized number of replicas ( $r_i/R$ ) for each of the items. To make the square-root trend clearer, for this particular plot, we have assumed query rates proportional to the id of the item, so the  $x$ -axis could also be named ‘query rate’. The plot includes the optimal square-root distribution (SR), drawn with a thick line. We have also experimented with other query rates, including Zipf-like ones, and the results were identical.

It should be clear from Fig. 1, that owner replication is far from the optimum. Path replication is better, but does not result in SR replication. Both PtP strategies, although different by nature, led to almost perfect SR replication. This also comes to confirm our intuition that the exact strategies used for the pull/push phases of PtP are not very important, as long as they are the same in both phases.

### 3. Consistency maintenance

Replication induces the need for consistency maintenance, that is, keeping the replicas up to date whenever changes occur. For the discussion that follows, we assume that each data item has a single *owner*, which is also the single peer that is allowed to modify the item<sup>3</sup>. Upon modification, the replicas which have been spread over the network must be made consistent with the most recent version of the data item.

The problem of consistency maintenance appears in many contexts [9, 10]. In [2, 6], various strategies were proposed in the context of peer-to-peer systems. In general, updates of a data item are broadcast by the owner and/or are searched for by the peers that have the replicas of the items. Thus, solutions to the consistency maintenance problem utilize:

- owner-initiated update *push*, so that peers with replicas are communicated the

---

<sup>3</sup>Actually, as noted in [6], others could also be allowed to modify it, by first notifying the owner, and then letting the latter update the replicas.

update,

- replica holder-initiated *pull*, either when needed or periodically, so as to discover new updates, if any, or
- a combination of the above.

It has been shown that usually a combined push/pull strategy (P/P for short) constitutes the best tradeoff between consistency levels and message overhead. The owner performs a limited push of the updates and the peers pull periodically, just in case the owner-initiated push did not reach them.

A basic problem in these P/P protocols is when should a peer pull. Pulling too often creates substantial message overhead. Pulling infrequently may result in missing important updates. Adaptive pull strategies try to minimize the communication overhead, while maintaining good consistency levels by having each replica holder pull at specific intervals. These intervals are determined by a time-to-refresh (TTR) parameter, which is adaptively adjusted depending on the previous pull results. If after the last pull the item was found unchanged, TTR is increased so as to pull less frequently; otherwise, TTR is decreased so as to check for updates more often.

Our premise is that efficient consistency maintenance can be achieved only in conjunction with efficient replication. That is, if the number of replicas and their placement is well-planned, then the algorithms for maintaining them under updates can be much more effective. To this end, we propose a novel push/pull update strategy that utilizes knowledge about replica creation so as to improve update efficiency.

Our experiments have shown that consistency maintenance can be achieved quite efficiently when replication is done in the optimal way, using the PtP strategies. Optimal replication not only minimizes the average search costs but also reduces the average update costs when combined with a suitable update strategy.

### 3.1. Updates under optimal replication

From now on we assume that items have been replicated in the network and that replication has been done using the PtP strategy. As discussed earlier, the PtP strategy requires that, after a successful search, the peer that found the item creates

a number of replicas, through a *replicate-push* phase, or R-push for short, with an appropriate TTL value. The basic idea now is to let this peer be held “responsible” for updating the replicas it created, as explained next. With respect to a particular data item, the nodes in the network fall into one of the following three categories:

- *owner*: the single peer that produces new versions of the data item
- *responsible*: a peer that searched for the item in the past (and thus forced the creation of replicas of the item)
- *indifferent*: a peer that was forced to hold a replica of the item.

The strategy, which we call PtPU, is a combination of push/pull. The owner broadcasts new updates to the network, through an update-push, or U-push for short. Whenever a “responsible” peer receives a new version of the item (either through an *update-pull* that it itself performed or an U-push that the item owner initiated), it undertakes the task of updating the replicas it created. In other words, it performs a U-push itself for the new version of the item. Moreover, this U-push should employ the same TTL parameter as the one used in the R-push, thereby reaching approximately the same nodes that were previously reached in order to create replicas.

This scheme has the potential of reducing the overhead of consistency maintenance significantly. A peer that is “responsible” for a resource should check (pull) frequently for newer updates of the item, using a smaller TTR value. Peers which were forced to have replicas of this item (“indifferent” peers) do not need to pull<sup>4</sup>, relying on some “responsible” peer to provide an update for them.

Summarizing, our strategy behaves as follows:

- The owner pushes the new versions of the item
- “Responsible” peers:
  - pull periodically,
  - push any updates they become aware of to their neighborhood exactly as when they created the replicas (i.e. with the same parameters as in the push phase of PtP).

---

<sup>4</sup>Or, they could pull quite infrequently (*cf.* the discussion in Section 4).

- The other (“indifferent”) peers do nothing; they rely on “responsible” peers to keep them updated.

For the periodic pulls of the “responsible” peers, we follow an adaptive scheme [6], whereby the time-to-pull-next (TTR) is decreased or increased according to the perceived version of the item. If the last pull did not return a newer version, the estimate for the next TTR will be increased by some constant:

$$\text{TTR}_e = \text{TTR} + C.$$

If, on the other hand, a more recent version of the item was found, the next TTR should be decreased. It should be decreased in proportion to the difference,  $D$ , in versions between the pulled item and the one the peer had — the higher the difference  $D$ , the more the missed updates, and hence the more frequent the pull (or, equivalently, the higher the decrease in TTR) should be. Thus, the estimate for the new TTR is:

$$\text{TTR}_e = \text{TTR}/(D + \beta),$$

where  $\beta$  is a parameter that provides some reduction in TTR in the case of  $D = 1$ .

The next TTR is a weighted average of the current TTR and the estimate, as follows:

$$\text{TTR} \leftarrow w\text{TTR}_e + (1 - w)\text{TTR},$$

where,  $w$  is a parameter determining the rate of change — smaller values of  $w$  make TTR change very slowly, while larger values make TTR adapt quickly to variations.

### 3.2. Experimental Evaluation

We have evaluated the performance of both the P/P and the PtPU strategies through extensive simulations. The network of peers is constructed and the data items are replicated using the PtP strategy as described in Section 2.2. After creating the replicas, we initiate simulation sessions. Each session runs for a number of rounds (turns). During each turn, the owner of an item creates a new version of the item with a given update probability  $p_u$  (update rate) and pushes it to the network. In the P/P strategy, all peers with replicas pull for new versions using adaptive pull. With

PtPU, only the “responsible” peers pull using, again, adaptive pull. In addition, the “responsible” peers push any received updates to their neighbors using exactly the same strategy used when the replicas were created (for example, using teeming with the same decay and TTL values).

We evaluate the performance of the update strategies with respect to two parameters: the achieved consistency and the associated message overhead. The consistency level is measured as the percentage of replicas that are up-to-date.

We experimented with different strategies for propagating the updates (i.e., pure flooding, random walkers, teeming and teeming with decay). The results attained were qualitative the same, thus, we report here only the results obtained when using teeming with decay, which is the method that gives us the most flexibility in terms of tuning the extend of the propagation. In particular, we present results when using three variations of teeming as summarized in Table 1. Wide teeming visits more peers, while narrow teeming produces smaller message overhead.

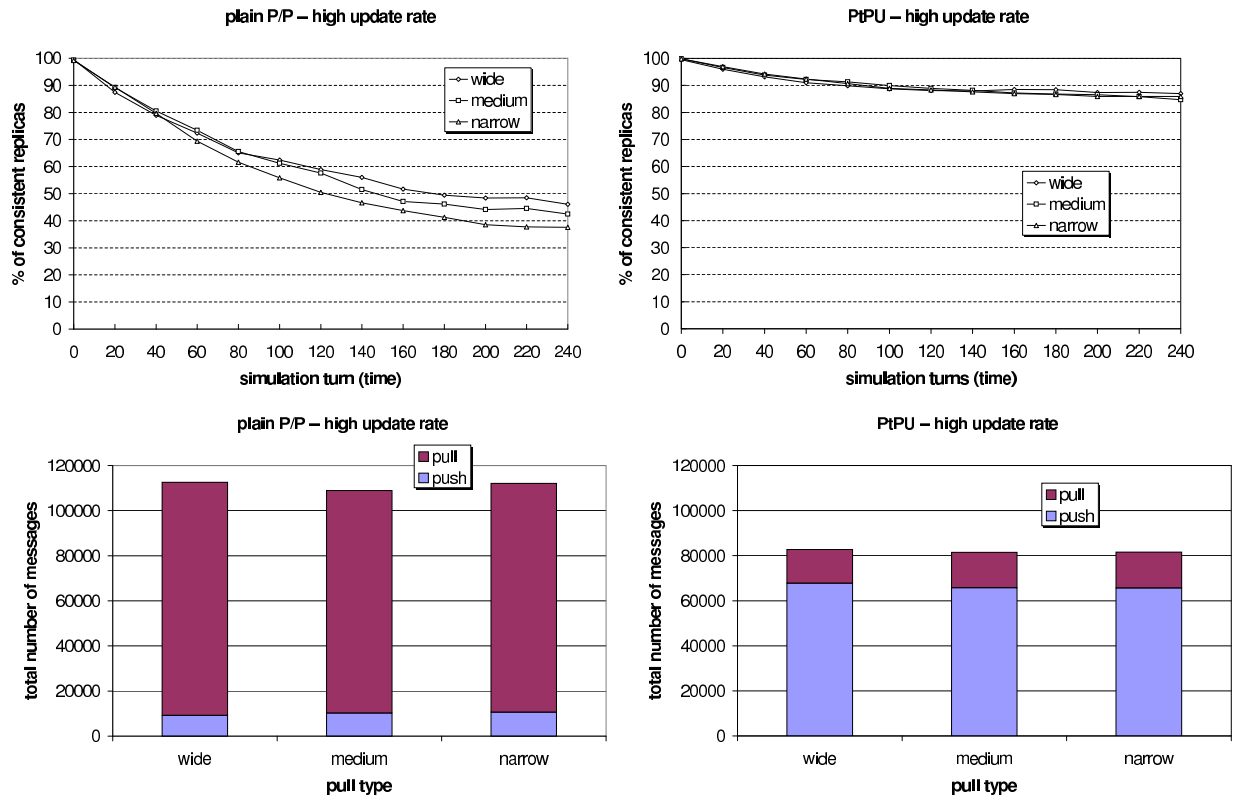
**Table 1.** *Versions of teeming*

<i>Extend of teeing</i>	<i>c (decay)</i>	<i>TTL</i>
Wide	0.1	5
Medium	0.3	5
Narrow	0.4	4

Regarding the adaptive pull, the tuning of its parameters is beyond the scope of this paper. A set of values that were found to work well in adapting the TTR is:  $w = 0.8$ ,  $b = 0.5$ , and  $C = 10$  turns, and those are the values that were used in all the experiments presented here. The reader is referred to [5, 9] for a detailed discussion of the topic.

### 3.2.1. Performance with respect to the update rate

The goal of the first set of experiments is to depict the behavior of plain P/P and PtPU under different update rates. We consider two cases: frequent updates ( $p_u = 0.1$ ), and infrequent updates ( $p_u = 0.025$ ).



**Figure 2.** Performance of the two strategies under high update rates.

The owner pushes the updates using narrow teeming. The reason for using such a rather limited push is for making the effect of pull more clear<sup>5</sup>. To discover a general trend, we let both strategies utilize exactly the same pull characteristics (i.e. the same variations of teeming) and see how they compare with each other.

The results are shown in Fig. 2 for high update rates and in Fig. 3 for infrequent updates. Each strategy is simulated for pulling with wide, medium and narrow teeming. In the case of high update rates, peers are forced to a high pull overhead in the P/P strategy so as to be frequently updated. In the PtPU case, though, pull is limited. Push messages are more since the “responsible” peers also propagate any updates they receive. For a low update rate, it is easier for any strategy to keep good consistency levels, utilizing fewer messages. Even in this case, though, PtPU achieved consistency levels above 92%, while plain P/P is, at best, a little

<sup>5</sup>A wide push would cause excessive message overhead and would keep most of the nodes updated, which in effect would make the pull phases useless.




Figure 3. Performance of the two strategies under low update rates.

above 80%. PtPU consistently outperforms P/P by any measure. It results in better consistency levels and, at the same time, with fewer messages.

### 3.2.2. Comparison of the two update policies

In this set of experiments, we compare further the two methods. In particular, we show (i) the level of consistency achieved when the two methods produce the same number of messages and (ii) the number of messages required by each method for achieving the same consistency level.

Here, we consider a medium update rate ( $p_u = 0.05$ ). For each strategy we repeatedly alter the pull parameters until we achieve the same value for the metric of interest (i.e. the consistency level or the number of messages) among all strategies.

The results are presented in Figs. 4–6. In the plots, we also consider the performance of P/P and PtPU, for the case where the creation of replicas does not follow

the PtP strategy. Instead, after the replication phase, the replicas get scattered across the network. Our goal is to show that losing the locality induced by the PtP strategy results in worsening the performance of both the P/P and PtPU strategies. Note that the number of replicas is kept the same; what differs is their placement in the network. The strategies under random placement of the replicas are marked with an “(R)” in the plots.

In Fig. 4 the owner uses a narrow push to propagate the updates. We run the simulator tuning the pull parameters until all strategies achieved approximately the same consistency level of 82%. The resulting message counts show that plain P/P required 43% more messages than PtPU to achieve the same consistency.

In Fig. 5 the owner uses a medium push to propagate the updates, so as to make it easier for the inferior strategies to achieve higher consistency levels (but, of course, with higher message overhead). The achieved consistency levels were approximately 95%. Once again, plain P/P required 46% more messages than PtPU.

In Fig. 6 all strategies generated approximately 62000 messages. PtPU required a narrow pull while P/P’s adaptive pull resulted in a wider teeming. The consistency levels observed show vividly the superiority of the PtPU strategy, which managed to achieve more than 90% consistency.

Another general conclusion from these plots is that, indeed, the random placement of replicas makes the performance of P/P and PtPU worse. This validates our intuition that the inherent locality of replica creation through PtP results in more efficient updates.

## 4. Discussion

In this paper, we consider replication in unstructured peer-to-peer systems. The idea behind our approach is that developing protocols for consistency maintenance which utilize knowledge about the strategy used to create the copies increases the efficiency of such protocols. Based on this, we develop a simple strategy for achieving square-root replication, which was previously proved to be optimal for unstructured peer-to-peer systems, and a consistency maintenance protocol that is tuned for our replication strategy.



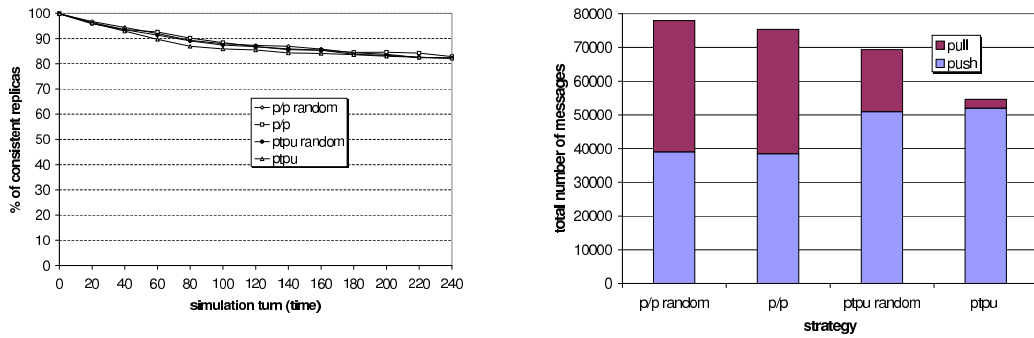


Figure 4. Comparison of the number of messages when all strategies result in consistency levels of approximately 82%.




Figure 5. Comparison of the number of messages when all strategies result in consistency levels of approximately 95%.

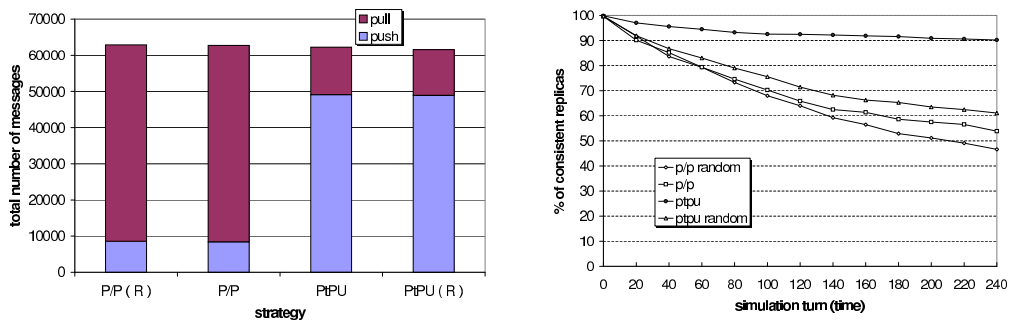


Figure 6. Comparison of consistency quality when all strategies generate the same number of messages.

Our experimental results show that our protocols achieve significantly better consistency for a smaller communication cost than protocols that do not exploit knowledge of the underlying replication strategy.

In our experiments we have assumed that the network does not change during the replication and update phases. We are currently studying the behavior of our strategies in more dynamic settings where peers enter or leave the system at will. In such an environment the PtPU strategy may encounter the following problem: a “responsible” peer could depart from the network, leaving thus a number of “indifferent” nodes without anybody to update their replicas for them. Thus, it is almost imperative that “indifferent” peers should pull, too, just in case the “responsible” node is not near them anymore.

## References

- [1] E. Cohen and S. Shenker, “Replication Strategies in Unstructured Peer-to-Peer Networks,” in *Proc. ACM SIGCOMM’02*, Aug. 2002.
- [2] A. Datta, M. Hauswirth, and K. Aberer, “Updates in highly unreliable, replicated peer-to-peer systems,” in *Proc. of ICDCS 2003, 23rd Int’l Conference on Distributed Computing Systems*, Providence, Rhode Island, May 2003, pp. 76–85.
- [3] V. V. Dimakopoulos and E. Pitoura, “Performance analysis of distributed search in open agent system,” in *Proc. IPDPS ’03, International Parallel and Distributed Processing Symposium*, Nice, France, May 2003.
- [4] V. Kalogeraki, D. Gunopoulos, and D. Zeinalipour-Yazti, “A local search mechanism for peer-to-peer networks,” in *Proc. CIKM ’02, 11th ACM Int’l Conference on Information and Knowledge Management*, McLean, Virginia, USA, Nov. 2002, pp. 300–307.
- [5] J. Lan, “Cache Consistency Techniques for Peer-to-Peer File Sharing,” MSc Thesis, Dept. of Computer Science, Univ. of Massachusetts, Tech. Rep., June 2002.
- [6] J. Lan, X. Liu, P. Shenoy, and K. Ramamritham, “Consistency maintenance in peer-to-peer file sharing networks,” in *Proc. of WIAPP’03, 3rd IEEE Workshop on Internet Applications*, San Jose, CA, USA, June 2003, pp. 76–85.
- [7] E. Leontiadis, V. V. Dimakopoulos, and E. Pitoura, “Cache Updates in a Peer-

- to-Peer Network of Mobile Agents,” in *Proc. P2P2005, 4th Int’l Conference of Peer to Peer Computing*, Zurich, Switzerland, Aug. 2004, pp. 10–17.
- [8] D. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker, “Search and Replication in Unstructured Peer-to-Peer Networks,” in *Proc. ICS’02, 16th ACM International Conference on Supercomputing*, New York, USA, June 2002.
- [9] R. Srinivasan, C. Liang, and K. Ramamritham, “Maintaining temporal coherency of virtual data warehouses,” in *Proc. RTSS ’98, 19th IEEE Real Time Systems Symposium*, Madrid, Spain, Dec. 1998.
- [10] B. Urgaonkar, A. Ninan, M. Raunak, R. Shenoy, and K. Ramamritham, “Maintaining mutual consistency for cached web objects,” in *Proc. ICDCS 2001, 21st Int’l Conference Distributed Computing Systems*, Phoenix, AZ, USA, Apr. 2001.