# Adaptive Push-Pull: Disseminating Dynamic Web Data

**Pavan Deolasee** **Amol Katkar** **Ankur Panchbudhe**
**Krithi Ramamritham** **Prashant Shenoy**

Department of Computer Science and Engineering.
Indian Institute of Technology Bombay
Mumbai, India 400076
{*pavan,amolk,ankurp,krithi*}@cse.iitb.ernet.in

Department of Computer Science
University of Massachusetts
Amherst, MA 01003
{*krithi,shenoy*}@cs.umass.edu

**Abstract**

An important issue in the dissemination of time-varying web data such as sports scores and stock prices is the maintenance of *temporal coherency*. In the case of servers adhering to the HTTP protocol, clients need to frequently *pull* the data based on the dynamics of the data and a user's coherency requirements. In contrast, servers that possess *push* capability maintain state information pertaining to clients and push only those changes that are of interest to a user. These two canonical techniques have complementary properties with respect to the level of temporal coherency maintained, communication overheads, state space overheads, and loss of coherency due to (server) failures. In this paper, we show how to combine push- and pull-based techniques to achieve the best features of both approaches. Our combined technique tailors the dissemination of data from servers to clients based on (i) the capabilities and load at servers and proxies, and (ii) clients' coherency requirements. Our experimental results demonstrate that such adaptive data dissemination is essential to meet diverse temporal coherency requirements, to be resilient to failures, and for the efficient and scalable utilization of server and network resources.

**Keywords**

Dynamic Data, Temporal Coherency, Scalability, Resiliency, World Wide Web, Data Dissemination, Push, Pull

## I. Introduction

Recent studies have shown that an increasing fraction of the data on the world wide web is time-varying (i.e., changes frequently). Examples of such data include sports information, news, and financial information such as stock prices. Web proxy caches that are deployed to improve user response times must track such dynamically changing data so as to provide users with temporally coherent information. The coherency requirements on a data item depends on the nature of the item and user tolerances. To illustrate, a user may be willing to receive sports and news information that may be out-of-sync by a few minutes with respect to the server, but may desire stronger coherency requirements on data items such as stock prices. A proxy can exploit user-specified coherency requirements by fetching and disseminating only those changes that are of interest and ignoring intermediate changes. For instance, a user who is interested in changes of more than a dollar for a particular stock price need not be notified of smaller intermediate changes.

In the rest of this section, we (a) describe the problem of temporal coherency maintenance in detail, (b) show the need to go beyond the canonical *Push*- and *Pull*-based data dissemination, and (c) outline the key

contributions of this paper, namely, the development and evaluation of adaptive protocols for disseminating dynamic, i.e., time-varying data.

### A. The Problem of Maintaining Temporal Coherency

Consider a proxy that caches several time-varying data items. To maintain coherency of the cached data, each cached item must be periodically refreshed with the copy at the server. We assume that a user specifies a temporal coherency requirement $tcr$ for each cached item of interest. The value of $tcr$ denotes the maximum permissible deviation of the cached value from the value at the server and thus constitutes the user-specified tolerance. Observe that $tcr$ can be specified in units of *time* (e.g., the item should never be out-of-sync by more than 5 minutes) or *value* (e.g., the stock price should never be out-of-sync by more than a dollar). In this paper, we only consider temporal coherency requirements specified in terms of the value of the object (maintaining temporal coherency specified in units of time is a simpler problem that requires less sophisticated techniques). As shown in figure 1, let $S(t)$, $C(t)$ and $U(t)$ denote the value of the data item at the server, cache and the user, respectively. Then, to maintain temporal coherency ($tc$) we should have

$$|U(t) - S(t)| \leq c.$$

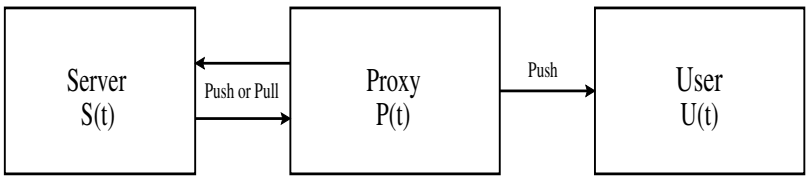


Fig. 1. The Problem of Temporal Coherency

The *fidelity* of the data seen by users depends on the degree to which their coherency needs are met. We define the fidelity $f$ observed by a user to be the total length of time that the above inequality holds (normalized by the total length of the observations). In addition to specifying the coherency requirement $tcr$, users can also specify their fidelity requirement $f$ for each data item so that an algorithm that is capable of handling users' fidelity and temporal coherency requirements ($tcr$s) can adapt to users' needs.

In this paper we develop *adaptive* push- and pull-based data dissemination techniques, to be employed between a server and a proxy, that maintain user-specified coherency and fidelity requirements. Techniques for disseminating data from proxies to end-users are not considered here, since resources such as network bandwidth are often plentiful on the proxy-user data path. Since proxies act as immediate clients to servers, henceforth, we use the terms proxy and client interchangeably (unless specified otherwise, the latter term is distinct from the ultimate end-users of data).

### B. The Need for Combining Push and Pull to Disseminate Dynamic Data

In the case of servers adhering to the HTTP protocol, proxies need to periodically *pull* the data based on the dynamics of the data and a user's coherency requirements. In contrast, servers that posses *push* capability maintain state information pertaining to clients and push only those changes that are of interest to a user/proxy.

The first contribution of this paper is an extensive evaluation of the canonical push and pull-based techniques using traces of real-world dynamic data. Our results, reported in Section II-C and summarized in Table I, show that these two canonical techniques have complementary properties with respect to resiliency to (server) failures, the level of temporal coherency maintained, communication overheads, state space overheads, and computation overheads.

TABLE I
BEHAVIORAL CHARACTERISTICS OF DATA DISSEMINATION ALGORITHMS

| **Algorithm** | **Resiliency** | **Temporal Coherency** (fidelity) | **Overheads (Scalability)** | | |
|---|---|---|---|---|---|
| | | | *Communication* | *Computation* | *State Space* |
| Push | Low | High | Low | High | High |
| Pull | High | Low (for small $tcr$s)<br>High (for large $tcr$s) | High | Low | Low |
| PaP | Graceful degradation | Adjustable (fine grain) | Low/Medium | Low/Medium | High |
| PoP | Delayed graceful degradation | Adjustable (coarse grain) | Low/Medium | Low/Medium | Low/Medium |
| PoPoPaP | Graceful degradation | Adjustable | Low/Medium | Low/Medium | Low/Medium |

Specifically, our results indicate that

- A pull-based approach does not offer high fidelity when the data changes rapidly or when the coherency requirements are stringent. Moreover, the pull-based approach imposes a large communication overhead (in terms of the number of messages exchanged) when the number of clients is large.
- A push-based algorithm can offer high fidelity for rapidly changing data and/or stringent coherency requirements. However, it incurs a significant computational and state-space overhead resulting from a large number of open push connections. Moreover, the approach is less resilient to failures due to its stateful nature.

These properties indicate that a push-based approach is suitable when a client requires its coherency requirements to be satisfied with a high fidelity, or when the communication overheads are the bottleneck. A pull-based approach is better suited to less frequently changing data or for less stringent coherency requirements, and when resilience to failures is important.

### C. Goal of the Proposed Algorithms

$T_{cr}$s of clients may vary across clients and bandwidth availability may vary with time, so a static solution to the problem of disseminating dynamic, i.e., time-varying, data will not be responsive to client needs or load/bandwidth changes. We need a *intelligent* and *adaptive* approach that can be tuned according to the client requirements and conditions prevailing in the network or at the server/proxy. Moreover, the approach should not sacrifice the scalability of the server (under load) or reduce the resiliency of the system to failures. One solution to this problem is to combine push and pull-based dissemination so as to realize the best features of both approaches while avoiding their disadvantages. The goal of this paper therefore is to develop techniques that combine push and pull in an *intelligent* and *adaptive* manner while offering good resiliency and scalability.

### D. Research Contributions of this Paper

In this paper, we propose two different techniques for combining push and pull-based dissemination.

1. $PaP$, our first algorithm, presented in Section III, simultaneously employs both push and pull to disseminate data, but has tunable parameters to determine the degree to which push and pull are used. Conceptually, the proxy is primarily responsible for pulling changes to the data; the server is allowed to push additional updates that are undetected by the proxy. By appropriate tuning, our algorithm can be made to behave as a push algorithm, a pull algorithm or a combination. Since both push and pull are simultaneously employed, albeit to different degrees, we refer to this algorithm as *Push-and-Pull (PaP)*.
2. $PoP$, our second algorithm, presented in Section IV, allows a server to adaptively choose between push- and pull-based dissemination for each connection. Moreover, the algorithm can switch each connection from push to pull and vice versa depending on the rate of change of data, the temporal coherency requirements and resource availability. Since the algorithm dynamically makes a choice of push or pull, we refer to it as *Push-or-Pull (PoP)*.

We have implemented our algorithms into a prototype server and a proxy. We demonstrate the efficacy of our approaches via simulations and an experimental evaluation. Complete source code for our prototype implementation and the simulator as well as the data used in our experiments is available from our web site.[1]

Table I summarizes the properties of our *PaP* and *PoP* algorithms vis-a-vis the canonical push and pull approaches.

The semantics of most of the entries in the table are self-evident even though the reason behind the stated properties of PaP and PoP will be clear only after they are described and evaluated. But a few words of explanation are in order.

[1] See http://www.cse.iitb.ernet.in/~krithi/dynamic.html.

- With respect to resiliency, PaP offers graceful degradation upon loss of state at the server or when the server loses a push connection. This is because, with PaP, a client normally obtains data through pushes *and* pulls, and when pushes from the server stop, pulls come to its rescue. So PaP seamlessly recovers from such failures. PoP is designed so that a client comes to know of state space losses or connection losses after a delay, at which point it needs to explicitly switch to pulling. Hence it too experiences graceful degradation, albeit after a delay. So, both PaP and PoP offer better failure handling properties than Push.
- The behavior of PaP and PoP can be adjusted to suit the temporal coherency requirements imposed on data. In the case of PaP, this is done by adjusting its parameters which can be done even on short time scales; with PoP, switching from Push to Pull or vice versa for a particular connection is viable over large time scales and this will change the temporal coherency of the disseminated data.
- The scalability properties of PoP and PaP are preferable to those of Pull or Push by themselves.

The last row of Table I shows the behavior of a protocol *PoPoPaP* that chooses one of Push, Pull, or PaP, thereby getting the benefits of all three where it is most appropriate to deploy them. This allows it to behave the best along all dimensions: resiliency, temporal coherency, and scalability.

## II. Push vs. Pull: Algorithms and their Performance

In this section, we present a comparison of push and pull-based data dissemination and evaluate their tradeoffs. These techniques will form the basis for our combined push-pull algorithms.

### A. Pull

To achieve $tc$ using a pull-based approach, a proxy can compute a *Time To Refresh (TTR)* attribute with each cached data item. The *TTR* denotes the next time at which the proxy should poll the server so as to refresh the data item if it has changed in the interim. A proxy can compute the *TTR* values based on the rate of change of the data and the user's coherency requirements. Rapidly changing data items and/or stringent coherency requirements result in a smaller TTR, whereas infrequent changes or less stringent coherency requirement require less frequent polls to the server, and hence, a larger *TTR*.[2] Observe that a proxy need not pull every single change to the data item, only those changes that are of interest to the user need to be pulled from the server (and the TTR is computed accordingly).

Clearly, the success of the pull-based technique hinges on the accurate estimation of the TTR value. Next, we summarize a set of techniques for computing the *TTR* value that have their origins in [22]. Given a user's coherency requirement, these techniques allow a proxy to adaptively vary the TTR value based on the rate

[2]Note that the *Time To Refresh (TTR)* value is different from the *Time to Live (TTL)* value associated with each HTTP request. The former is computed by a proxy to determine the next time it should poll the server based on the $tcr$; the latter is provided by a web server as an estimate of the next time the data will be modified.

of change of the data item. The TTR decreases dynamically when a data item starts changing rapidly and increases when a hot data item becomes cold. To achieve this objective, the *Adaptive TTR* approach takes into account (a) static bounds so that TTR values are not set too high or too low, (b) the most rapid changes that have occurred so far and (c) most recent changes to the polled data.

In what follows, we use $D_0, D_1, \ldots, D_l$ to denote the values of a data item $D$ at the server in chronological order. Thus, $D_l$ is the most recent value data item $D$.

The adaptive *TTR* is computed as:
$TTR_{adaptive} = Max(TTR_{min},\ Min(TTR_{max},\ a \times TTR_{hr} + (1-a) \times TTR_{dyn}))$
where

- $[TTR_{min}, TTR_{max}]$ denote the range within which TTR values are bound.
- $TTR_{hr}$ denotes the most conservative, i.e., smallest, TTR value used so far. If the next TTR is set to $TTR_{hr}$, temporal coherency will be maintained even if the maximum rate of change observed so far recurs. However, this TTR is pessimistic since it is based on worst case rate of change at the source. If this worst case rapid change occur for only a small duration of time, then this approach is likely to waste a lot of bandwidth especially if the user can handle some loss of fidelity.
- $TTR_{dyn}$ is a learning based TTR estimate founded on the assumption that the dynamics of the last few (two, in the case of the formula below) recent changes are likely to be reflective of changes in the near future.

$$TTR_{dyn} = (w \times TTR_{estimate}) + ((1-w) \times TTR_{latest})$$

where

– $TTR_{estimate}$ is an estimate of the TTR value, based on the most recent change to the data.

$$TTR_{estimate} = \frac{TTR_{latest}}{|D_{latest} - D_{penultimate}|} \times c$$

If the recent rate of change persists, $TTR_{estimate}$ will ensure that changes which are greater than or equal to $c$ are not missed.

– weight $w$ ($0.5 \leq w < 1$, initially 0.5) is a measure of the relative change between the recent and the old changes, and is adjusted by the system so that we have the *recency* effect, i.e., more recent changes affect the new $TTR$ more than the older changes.

- $0 \leq a \leq 1$ is a parameter of the algorithm and can be adjusted dynamically depending on the fidelity desired, with a higher fidelity demanding a higher value of $a$.

The adaptive TTR approach has been experimentally shown to have the best $tc$ properties among several TTR assignment approaches [22]. Consequently, we choose this technique as the basis for pull-based dissemination.

### B. Push

In a push-based approach, the proxy registers with a server, identifying the data of interest and the associated $tcr$, i.e., the value $c$. Whenever the value of the data changes, the server uses the $tcr$ value $c$ to determine if the new value should be pushed to the proxy; only those changes that are of interest to the user (based on the $tcr$) are actually pushed. Formally, if $D_k$ was the last value that was pushed to the proxy, then the current value $D_l$ is pushed if and only if $|D_l - D_k| \geq c, 0 \leq k \leq l$. To achieve this objective, the server needs to maintain state information consisting of a list of proxies interested in each data item, the $tcr$ of each proxy and the last update sent to each proxy.

The key advantage of the push-based approach is that it can meet stringent coherency requirements—since the server is aware of every change, it can precisely determine which changes to push and when.

But, an important consideration in push servers is the amount of state required to be maintained at the server (Section II-C). To achieve a reduction in the space needed, rather than maintain the data and $tcr$ needs of individual client separately, the server combines all requests for a particular data item $D$ and needing a particular $tcr$; as soon as the change to $D$ is greater than equal to $c$, all the clients associated with $D$ are notified.

Let the above optimization process convert $n$ connections into $u$ unique $(D, c)$ pairs. Now for one unique $(D, c)$,

$$u \times (bytes\ needed\ for\ a\ (D,c)\ pair)\ +\ n \times (bytes\ needed\ for\ a\ connection\ state) \tag{1}$$

Also, since $u \leq n$, this space is less than the space required if above optimization was not applied (in which case $u$ in first term of 1 will be replaced by $n$).

But even with this optimization, the amount of state at the server can be considerable and a loss of this state can drastically degrade the resiliency and scalability of the push server. Hence it is important to consider techniques that deliver the advantages of push and pull respectively without suffering from the disadvantages of either. We propose one such technique in section IV.

### C. Performance of Push vs. Pull

In what follows, we compare the push and pull approaches along several dimensions: maintenance of temporal coherency, communication overheads, computational overheads, space overheads, and resiliency.

#### C.1 Experimental Model

These algorithms were evaluated using a prototype server/proxy that employed trace replay. For Pull, we used a vanilla HTTP web server with our prototype proxy. For Push, we used a prototype server that

uses unicast and connection-oriented sockets to push data to proxies. All experiments were done on a local intranet. We also ran carefully instrumented experiments on the internet and the trends observed were consistent with our results.

Note that it is possible to use multicast for push; however, we assumed that unicast communication is used to push data to each client (thus, results for push are conservative upper-bounds; the message overheads will be lower if multicast is used).

### C.2 Traces Used

Quantitative performance characteristics are evaluated using real world stock price streams as exemplars of dynamic data. The presented results are based on stock price traces (i.e., history of stock prices) of a few companies obtained from `http://finance.yahoo.com`. The traces were collected at a rate of 2 or 3 stock quotes per second. Since rate of change of any stock quote is much greater than even 1 change per second, the traces can be considered to be "real-time" traces. For empirical and repetitive evaluations, we "cut out" the history for the time intervals listed in table II and experimented with the different mechanisms by determining the stock prices they would have observed had the source been live. A trace that is 2 hours long, has approximately 15000 data values. All curves portray the averages of the plotted metric over all these traces. Few of the experiments were done with quotes obtained in real-time, but the difference was found to be negligible when compared to the results with the traces.

The Pull approach was evaluated using the Adaptive TTR algorithm with an $a$ value of 0.9, $TTR_{min}$ of 1 second and three $TTR_{max}$ values of 10, 30 and 60 seconds.

TABLE II
TRACES USED FOR THE EXPERIMENT

| **Company** | **Date** | **Time** |
|---|---|---|
| Dell | Jun 1, 2000 | 21:56-22:53 IST |
| UTSI | Jun 1, 2000 | 22:41-23:15 IST |
| CBUK | Jun 2, 2000 | 18:31-21:57 IST |
| Intel | Jun 2, 2000 | 22:14-01:42 IST |
| Cisco | Jun 6, 2000 | 18:48-22:20 IST |
| Oracle | Jun 7, 2000 | 00:01-01:59 IST |
| Veritas | Jun 8, 2000 | 21:20-23:48 IST |
| Microsoft | Jun 8, 2000 | 21:02-23:48 IST |

### C.3 Maintenance of Temporal Coherency

Since a push-based server communicates every change of interest to a connected client, a client's $tcr$ is never violated as long as the server does not fail or is so overloaded that the pushes are delayed. Thus, a push-based server is well suited to achieve a fidelity value of 1. On the other hand, in the case of a pull-based server, the frequency of the pulls (translated in our case to the assignment of TTR values) determines the degree to which client needs are met. We quantify the achievable fidelity of pull-based approaches in terms of the probability that user's $tcr$ will be met. To do so, we measure the durations when $|U(t) - S(t)| > c$. Let $\delta_1, \delta_2, \ldots, \delta_n$ denote these durations when user's $tcr$ is violated. Let $observation_interval$ denote the total time for which data was observed by a user. Then fidelity is

$$1 - \frac{\sum_{i=1}^{n} \delta_i}{observation_interval}$$

and is expressed as a percentage. This then indicates the percentage of time when a user's desire to be within $c$ units of the source is met.

Figure 2-(a) shows the fidelity for a pull-based algorithm that employs adaptive *TTRs* . Recall that the Push algorithm offers a fidelity of 100%. In contrast, the figure shows that the *pull* algorithm has a fidelity of 70-80% for stringent coherency requirements and its fidelity improves as the coherency requirements become less stringent. (The curve marked PaP is for the *PaP* algorithm that combines Push and Pull and is described in Section III-A).

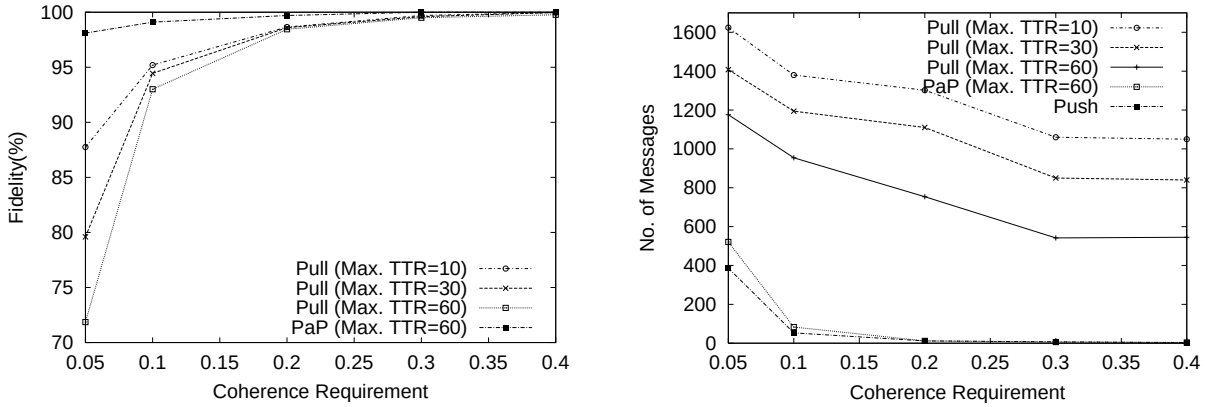


(a) Fidelity (b) Overheads

Fig. 2. Fidelity and Overheads for Varying Coherence Requirements

### C.4 Communication Overheads

In a push-based approach, the number of messages transferred over the network is equal to the number of times the user is informed of data changes so that the user specified temporal coherency is maintained. (In a network that supports multicasting, a single push message may be able to serve multiple clients.)

A pull-based approach requires two messages—an HTTP IMS request, followed by a response—per poll. Moreover, in the pull approach, a proxy polls the server based on its estimate of how frequently the data is changing. If the data actually changes at a slower rate, then the proxy might poll more frequently than necessary. Hence a pull-based approach is liable to impose a larger load on the network. However, a push-based approach may push to clients who are no longer interested in a piece of information, thereby incurring unnecessary message overheads. We quantify communication overheads in terms of the number of messages exchanged between server and proxy. Figure 2-(b) shows the variation in the number of messages with coherence requirement $0.05 \leq c \leq 0.4$. As seen in figure 2-(b), the *Push* approach incurs a small communication overhead because only values of interest to a client are transferred over the network. The *Pull* approach, on the other hand, imposes a significantly higher overhead.

### C.5 Computational Overheads

Computational overheads for a pull-based server results from the need to deal with individual pull requests. After getting a pull request from the proxy, the server has to just look up the latest data value and respond. On the other hand, when the server has to push changes to the proxy, for each change that occurs, the server has to check if the $tcr$ for any of the proxies has been violated. This computation is directly proportional to the rate of arrival of new data values and the number of unique temporal coherency requirements associated with that data value. Then the computational overhead per data item is of the order of rate of arrival of new values times the number of unique coherence requirements associated with that a value. Although this is a time varying quantity in the sense that the rate of arrival of data values as well as number of connections change with time, it is easy to see that push is computationally more demanding than pull. On the other hand, it is important to remember that servers respond to individual pull requests and so may incur queueing related overheads.

### C.6 Space Overheads

A pull-based server is stateless. In contrast, a push-based server must maintain the $tcr$ for each client, the latest pushed value, along with the state associated with an open connection. Since this state is maintained throughout the duration of client connectivity, the number of clients which the server can handle may be limited when the state space overhead becomes large (resulting in scalability problems). That is, the push based approach can suffer from scalability problems arising from the need to store large amounts of state information.

From the above discussion we can see that:

- Computational overheads at a server are directly proportional to the sum of the number of data items requested by clients.

- State space overheads at a server are directly proportional to the number of push connections and the number of data items served (also see equation 1).

### C.7 Resiliency

By virtue of being stateless, a pull-based server is resilient to failures. In contrast, a push-based server maintains crucial state information about the needs of its clients; this state is lost when the server fails. Consequently, the client's coherency requirements will not be met until the proxy detects the failure and re-registers the $tcr$ requirements with the server.

The above results are summarized in Table I. In what follows, we present two approaches that strive to achieve the benefits of the two complementary approaches by adaptively combining Push and Pull.

## III. PaP: Dynamic Algorithm with *Push and Pull* Capabilities

In this section, we present *Push-and-Pull (PaP)* — a new algorithm that simultaneously employs both push and pull to achieve the advantages of both approaches. The algorithm has tunable parameters that determine the degree to which push and pull are employed and allow the algorithm to span the entire range from a push approach to a pull approach. Our algorithm is motivated by the following observations.

The pull-based adaptive TTR algorithm described in Section II-A can react to variations in the rate of change of a data item. When a data item starts changing more rapidly, the algorithm uses smaller TTRs (resulting in more frequent polls). Similarly, the changes are slow, TTR values tend to get larger. If the algorithm detects a violation in the coherency requirement (i.e., $|D_{latest} - D_{penultimate}| > c$), then it responds by using a smaller TTR for the next pull. A further violation will reduce the TTR even further. Thus, successive violations indicate that the data item is changing rapidly and the proxy gradually decreases the TTR until the TTR becomes sufficiently small to keep up with the rapid changes. Experiments reported in [22] show that the algorithm gradually "learns" about such "clubbed" (i.e, successive) violations and reacts appropriately. So, what we need is a way to prevent even the small number of $tc$ violations that occur due to the delay in this gradual learning process. Furthermore, if a rapid change occurs at the source and then the data goes back to its original value before the next pull, this "spike" will go undetected by a pull-based algorithm. The PaP approach described next helps the TTR algorithm to "catch" all the "clubbed" violations properly; moreover "spikes" also get detected. This is achieved by endowing push capabilities to servers and having the server push changes that a proxy is unable to detect. This increases the fidelity for clients at the cost of endowing push capability to servers. Note that, since proxies continue to have the the ability to pull, the approach is more resilient to failures than a push approach (which loses all state information on a failure).

### A. The PaP Algorithm

Suppose a client registers with a server and intimates its coherency requirement $tcr$. Assume that the client pulls data from the server using an algorithm, say $A$, to decide its TTR values (e.g., *Adaptive TTR*). After initial synchronization, server also runs algorithm $A$. Under this scenario, the server is aware of when the client will be pulling next. With this, whenever server sees that the client must be notified of a new data value, the server pushes the data value to the proxy if and only if it determines that the client will take time to poll next. The state maintained by this algorithm is a soft state in the sense that even if push connection is lost or the clients' state is lost due to server failure, the client will continue to be served at-least as well as under $A$. Thus, compared to a Push-based server, this strategy provides for graceful degradation.

In practice, we are likely to face problems of synchronization between server and client because of variable network delays. Also, the server will have the additional computational load imposed by the need to run the TTR algorithm for all the connections it has with its clients. The amount of additional state required to be maintained by the server cannot be ignored either. One could argue that we might as well resort to Push which will have the added advantage of reducing the number of messages on the network. However, we will have to be concerned with the effects of loss of state information or of connection loss on the maintenance of $tc$.

Fortunately, for the advantages of this technique to accrue, the server need not run the full-fledged TTR algorithm. A good approximation to computing the client's next TTR will suffice. For example, the server can compute the difference between the times of the last two pulls ($diff$) and assume that the next pull will occur after a similar $delay$, at $t_{predict}$. Suppose $T(i)$ is the time of the most recent value. The server computes $t_{predict}$, the next predicted pulling time as follows:

- let $diff = T(i) - T(i-1)$
- server predicts the next client polling time as $t_{predict} = T(i) + diff$.

If a new data value becomes available at the server before $t_{predict}$ and it needs to be sent to the client to meet the client's $tcr$, the server pushes the new data value to the client.

In practice, the server should allow the client to pull data if the changes of interest to the client occur close to the client's expected pulling time. So, the server waits, for a duration of $\epsilon$, a small quantity close to $TTR_{min}$, for the client to pull. If a client does not pull when server expects it to, the server extends the push duration by adding $(diff - \epsilon)$ to $t_{predict}$. It is obvious that if $\epsilon = 0$, *PaP* reduces to push approach; if $\epsilon$ is large then the approach works similar to a pull approach. Thus, the value of $\epsilon$ can be varied so that the number of pulls and pushes is balanced properly. $\epsilon$ is hence one of the factors which decides the $tc$ properties of the PaP algorithm as well as the number of messages sent over the network.

### B. Details of PaP

The arguments at the beginning of this section suggest that it is a good idea to let the proxy pull when it is polling frequently anyway and violations are occurring rapidly. Suppose, starting at $t_i$ a series of rapid changes occurs to data $D$. This can lead to a sequence of "clubbed" violations of $tcr$ unless steps are taken. The adaptive TTR algorithm triggers a decrease in the TTR value at the proxy. Let this TTR value be $TTR_k$. The proxy polls next at $t_{i+1} = t_i + TTR_k$. According to the PaP algorithm, the server pushes any data changes above $tcr$ during $(t_i, t_{i+1})$. Since a series of rapid changes occurs, the probability that some violation(s) may occur in $(t_i, t_{i+1}]$ is very high and thus these changes will also be pushed by the server further forcing a decrease in the TTR at the proxy and causing frequent polls from the proxy. Now, the TTR value at the proxy will tend towards $TTR_{min}$ and $diff$ will also approach zero, thus making the durations of possible pushes from the server close to zero. It is evident that if rapid changes occur, after a few pushes, the push interval will be zero, and client will pull almost all the rapid changes thereafter. Thus the server has helped the proxy pull sooner than it would otherwise. This leads to better fidelity of data at the proxy than with a pull approach.

If an isolated rapid change (i.e., spike) occurs, then the server will push it to the proxy leading to a decrease in the TTR used next by the proxy. It will poll sooner but will not find any more violation and that in turn will lead to an increase in the TTR.

Thus, the proxy will tend to pull nearly all but the first few in a series of rapid changes helped by the initial pushes from the server, while all "spikes" will be pushed by the server to the proxy. The result is that all violations will be caught by the PaP algorithm in the ideal case (e.g., with the server running the adaptive TTR algorithm in parallel with the proxy). In case the server is estimating the proxy's next TTR, the achieved $tc$ can be made to be as close to the ideal, as exemplified by Pure Push, by proper choice of $\epsilon$.

Overall, since the proxy uses the pushed (as well as pulled) information to determine TTR values, the adaptation of the TTRs would be much better than with a pull-based algorithm alone.

Although the amount of state maintained is nearly equal to push, the state is a soft state. This means that even if the state is lost due to some reason or the push connection with a proxy is lost, the performance will be at-least as good as that of TTR algorithm running at the proxy as clients will keep pulling.

### C. Performance Analysis of PaP

Figure 2-(a) shows that for *PaP* algorithm, the fidelity offered is more than 98% for stringent $tcr$ and 100% for less stringent $tcr$. Let us now reexamine figure 2-(b). Compared to Pull, the *PaP* algorithm has very little network overhead because of the push component. Its network overheads are, however, slightly higher than that of *Push*.

Figure 3-(a) shows the variation in the amount of push versus pull when we change $a$ (adjustment factor, see section II-A) for a fixed $c$ of 0.05. Large values of $a$ make pulls more successful and vice-versa. Another point to note is that the slopes for both the curves are very small, thus $\epsilon$ does not change the proportion of pushes and pulls considerably.

If we combine the trends shown in figure 2-(b), it is clear that pull is ineffective for stringent coherency requirements and it can meet less stringent requirements, albeit at a larger cost. Consider the data points, when c=0.05, the Pull approach incurs 200% more polls (messages) and yet provides only a fidelity of 70-80% (i.e., 20% temporal consistency violations). At c=0.1, the fidelity provided improves to 95% but cost of polls is still very high (factor of 10 larger). These trends indicate that Pull is suitable when the coherency requirements are less stringent and the data changes less frequently. The figures also show that PaP can provide near-perfect fidelity at a low message cost, always.

The value of $TTR_{max}$ needs to be chosen to balance the number of pushes and pulls. The results in figure 3-(b), as one would expect, show that when $TTR_{max}$ is large the number of successful pushes is large, but as we decrease $TTR_{max}$, the number of pushes decreases slowly until a point where pulls start dominating (point around $TTR_{max} = 10$ in the figure). After this point, pushes decrease while pulls increase rapidly.

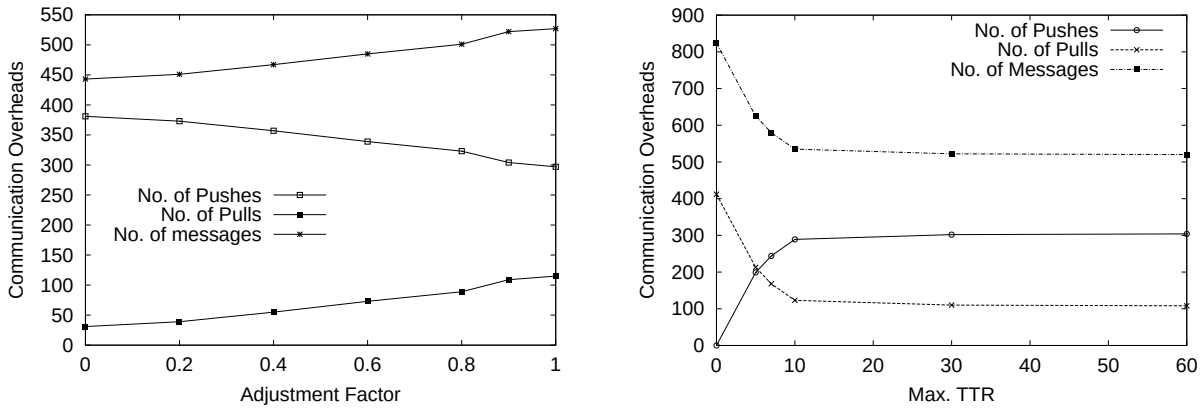


(a) Effect of Adjustment Factor on Overheads

(b) Effect of Max. TTR on Overheads

Fig. 3. Effect of Parameters on Overheads

Figure 4-(a) shows the variation in fidelity when $\epsilon$ is varied. When $\epsilon$ is zero, the algorithm reduces to push and hence fidelity is 100%. But as we start increasing the value of $\epsilon$ the fidelity starts suffering. For values of $\epsilon < 3$, the fidelity is above 75%. And for $\epsilon = TTR_{min} = 1$, fidelity is approximately 99%. For values of $\epsilon$ closer to $TTR_{max}$ (in this case 60), fidelity is low as the pulls overtake pushes and the algorithm behaves like a TTR algorithm.

Figure 4-(a) also shows the effect of changing $TTR_{max}$ in conjunction with $\epsilon$ on the fidelity offered by the algorithm. From 3-(b) we know that as $TTR_{max}$ decreases, pulls increase. As pulls become more dominant,

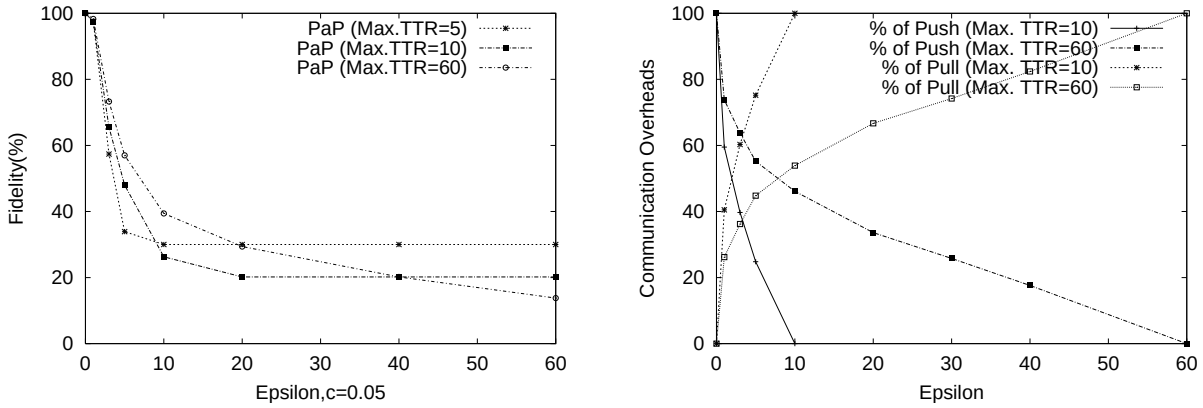


(a) Effect of epsilon on Fidelity

(b) Effect of epsilon on Overheads

Fig. 4. Effect of epsilon

the server has less chance to push the data values, and a bigger $\epsilon$ gives the server fewer opportunities to push. This explains the effect in 4-(a) for $TTR_{max} = 5$ or $TTR_{max} = 10$. As pulls increase and the server has less and less chance to push, fidelity suffers and decreases more rapidly than in the case of $TTR_{max} = 60$. It can also be observed that, as $\epsilon$ takes values greater than $TTR_{max}$, fidelity offered becomes constant. This is because even if server sets $\epsilon$ greater than $TTR_{max}$ client will keep polling at the maximum rate of $TTR_{max}$. In effect, setting $\epsilon$ greater than $TTR_{max}$ is equivalent to setting it to $TTR_{max}$. This explains the crossover of curves in figure 4-(a).

Figure 4-(b) also portrays the variation in the number of successful pushes and pulls. As expected, as $\epsilon$ is increased the number of pulls become higher and higher. For $\epsilon = 0$, there are no pulls and for $\epsilon = TTR_{max}(= 60)$ there are no pushes. If we compare the graphs in figure 4-(b), we can see that more fidelity requires more number of pushes and for the case where number of pushes is equal to number of pulls, fidelity is close to 50%. The more we increase the number of pulls (i.e., $\epsilon$), the lower the obtained fidelity.

The graphs for $TTR_{max} = 10$ in figure 4-(b) show a similar trend to those for $TTR_{max} = 60$, over the values $\epsilon = 0$ to $\epsilon = 10$ and 60 respectively. This is expected, as $\epsilon$ takes values between $TTR_{min}$ and $TTR_{max}$, the behavior of the algorithm changes from being like pure push to pure pull (i.e. percentage of push drops). Setting $\epsilon = TTR_{max}$ makes the percentage of push equal to zero.

### *D. Tuning of PaP: Making the Approach Adaptive and Intelligent*

One of the primary goals of our work was to have an adaptive and intelligent approach which can be tuned according to conditions prevailing in the system. From the above analysis of the *PaP* algorithm it is clear that it has a set of tunable parameters using which one can achieve Push capability or Pull capability or anything in between. So it is fairly flexible. The parameter that dominates the scenario the most is $TTR_{max}$ in terms of communication overheads, while the other parameter is $\epsilon$. If the user wants high fidelity then we must set

$\epsilon$ close to $TTR_{min}$. So, if a user is flexible with respect to fidelity, then $\epsilon$ can be set accordingly. But we also know that the rate of change in fidelity with $\epsilon$ depends on $TTR_{max}$. So, if both available bandwidth and fidelity can be adjusted, then we can set $TTR_{max}$ properly. Some of the typical scenarios are:

- if we know that the bandwidth available is low and still high fidelity is desired, then we choose a moderate $TTR_{max}$ (close to 10 as in figure 3-(b)) and $\epsilon$ low (close to 2 as in Figure 4-(a)).
- if the bandwidth is scarce and fidelity desired is also not high, then we can set $TTR_{max}$ and $\epsilon$ both to a moderate value (close to 15 and 5 respectively).
- if the bandwidth available is high and fidelity desired is also high, then we can set $TTR_{max}$ low (close to 5) and $\epsilon$ equal to $TTL_{min}$, thus having less pushes (and more pulls) but still good fidelity.
- If the $tcr$ is not stringent, then a lower value can be set for $TTR_{max}$, thereby making the system resort to pulls.
- if the load on the server is high (due to more push), $\epsilon$ can be set to a moderate/high value and/or $TTR_{max}$ can be set to low/moderate value so that the amount of push decreases and there are more pulls.

## IV. PoP: Dynamically Choosing Between *Push or Pull*

Whereas PaP achieves its adaptiveness through the adjustment of parameters such as $\epsilon$ and $TTR_{max}$ and thereby obtains a range of behaviors with push and pull at the two extremes, we now describe a somewhat simpler approach wherein, based on the availability of resources and the data and $tc$ needs of users, a server chooses push or pull for a particular client. Consequently, we refer to our approach as *Push-or-Pull (PoP)*.

### *A. The PoP Algorithm*

PoP is based on the premise that at any given time a server can categorize its clients either as push clients or pull clients and this categorization can change with system dynamics. This categorization is possible since the server knows the parameters like the number of connections it can handle at a time and can determine the resources it has to devote to each mode (Push/Pull) of data dissemination so as to satisfy its current clients. The basic ideas behind this approach are:

- allow failures at a server to be detected early so that, if possible, clients can switch to pulls, and thereby achieve graceful degradation to such failures. To achieve this, servers are designed to push data value to their push clients when one of two conditions is met:
  1. The data value at the server differs from the previously forwarded value by $tcr$ or more.
  2. A certain period of time $TTR_{limit}$ has passed since the last change was forwarded to the clients.

The first condition ensures that the client is never out of sync with the values at the server by an amount exceeding the $tcr$ of the client. The second condition assures the client after passage of every $TTR_{limit}$ interval that (a) the server is still up and (b) the state of the client with the server is not lost. This makes

the approach resilient. In case of the state of the client being lost or the connection being closed because of network errors, the client will come to know of the problem after $TTR_{limit}$ time interval, after which the client can either request the server to reinstate the state or start pulling the data itself. This ensures that in the worst case, the time for which the client remains out of sync with the server never exceeds $TTR_{limit}$.

- equipped with this, the server can be designed to provide push service as the default to all the clients provided it has sufficient resources.
- when a resource constrained situation arises (upon the registration of a new client or network bandwidth changes) some of the push-based clients are converted to become pull-based clients based on the criteria that we had determined earlier.

The state diagram for achieving this adaptation is shown in figure 5.

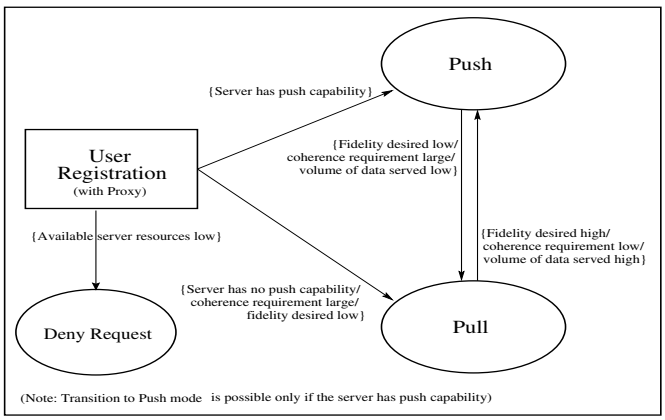


Fig. 5. PoP: Choosing between Push and Pull

### *B. Details of PoP*

Whenever a client contacts a server for a data item, the client also specifies its $tcr$ and fidelity requirements.

- Irrespective of the fidelity requirement, if the server has sufficient resources (such as a new monitoring thread, memory, etc.), the client is given a push connection.
- Otherwise, if the client can tolerate lower fidelity, then server disseminates data to that client based on client pull requests.
- If the request desires 100% fidelity and the server does not have sufficient resources to satisfy it, then the server takes steps to convert some push clients to pull. If this conversion is not possible, then the new request is denied.

In the latter case, the push clients chosen are those who can withstand the resulting degraded fidelity, i.e., those who had originally demanded less than 100% fidelity but had been offered higher fidelity because resources were available then for push connections. Which client(s) to choose is decided based on additional

considerations including (a) bandwidth available (b) rate of change of data and (c) $tcr$ . If bandwidth available with a client is low, then forcing the client to pull will only worsen its situation since pull requires more bandwidth than push. If the rate of change of data value is low or the $tcr$ is high, then pull will suffice. Thus, from amongst the clients which had specified low fidelity requirement, we choose proxies which have (a) specified a high value of $tcr$, or (b) volume of data served is small. If a suitable (set of) client(s) is found, the server sends appropriate "connection dropped" intimation to the client so that it can start pulling.

### *C. Performance of PoP*

Using the same traces given in table II we evaluated PoP. The experiments were performed by running the server on a load free Pentium-II machine and simulating clients from four different load free Pentium-II machines. There were 56 users on each client machine, accessing 3-4 data items. Keeping the resources constant, the ratio of push to pull connections was varied and the effect on average fidelity experienced by clients in pull mode as well as across all the clients was measured.

Experimental results plotted in figure 6-(a) indicate that the communication overheads drop when the percentage of push sockets is increased. This is to be expected because push algorithms are optimum in terms of communication overheads. As we increase the percentage of push sockets, while the push clients may be able to experience 100% coherence, the percentage of pull requests that are denied due to queue overflow grows exponentially. These results indicate that a balance between pull and push connections must be struck if we want to achieve scalability while achieving high coherency.

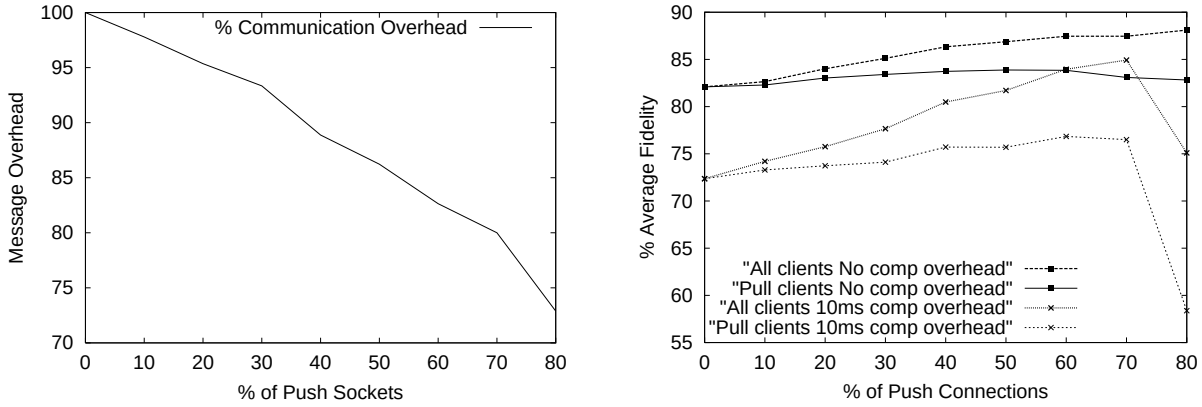


(a) Communication Overhead (b) Average Fidelity

Fig. 6. Variation with % of Push Connections

We measured the effect of increasing the percentage of push connections on fidelity. As the number of push connections increases, proxies which were serving the largest number of data items or data items with stringent temporal coherency requirements are moved from pull to push mode. The implemented system worked with a fixed number of clients/data items and so the results below do not reflect the effect of

admission control (i.e. request denial based on server load and client profile, which includes its requirements, volume of data served to it and its network status) that is part of PoP. The results are plotted in figure 6-(b) for two cases of computational overheads per pull request: (1) no computational overheads, except for those connected with use of the socket, and (2) a 10 msec computational overhead per pull request, in addition to the socket handling overheads.

When the computational overheads for a pull are negligible, average fidelity across all clients improves gradually as we increase the percentage of push clients. When a small computational overhead of 10 msecs per pull is added, while fidelity improves up to a point, when the number of pull connections becomes small, some of the pull requests experience denial of service thereby affecting the average fidelity across all clients. In fact, the overall fidelity drops nearly 10%.

Recall that all push clients experience 100% fidelity. So, the above drop in fidelity is all due to the pull clients. This is clear when we study the variation of the average fidelity of pull clients. With zero computational overheads for pulls, as we increase the number of push clients, fidelity for pull clients improves from 82% to to 84% before dropping to 83%. The improvement as well as drop is more pronounced under 10 msec pull overheads. When a large percentage of the clients are in pull mode, the number of pull requests is very high. This increases the average response time for each client, which in effect, decreases the fidelity for pull clients. This scalability problem is due to computation load at the server when a large number of pull clients are present. As more and more clients switch to push mode, the number of pull requests drops, the response time of the server improves, and better fidelity results. The fidelity for pull clients peaks and then starts deteriorating. At this point the incoming requests cause overflows in the socket queues and the corresponding requests are denied. These again cause an increase in the effective response time of the client and fidelity decreases. The last portion of the curve clearly brings out the scalability issue arising because of resource constraints.

These results clearly identify the need for the system to allocate push and pull connections intelligently. An appropriate allocation of push and pull connections to the registered clients will provide the temporal coherency and fidelity desired by them. In addition, when clients request access to the server and the requisite server resources are unavailable to meet needs of the client, access must be denied. As figure 5 indicates, this is precisely what PoP is designed to do.

### *D. Tuning of PoP*

It is clear from the results plotted in figure 6-(a) and 6-(b), that the way in which clients are categorized as push and pull clients affects the performance of PoP. So the system must dynamically and intelligently allocate the available resources amongst push and pull clients. For example,

- when the system has to scale to accommodate more clients, it should preempt a few push clients (ideally,

those which can tolerate some loss in fidelity) and give those resources to pull clients.

- if the number of clients accessing the server is very small, a server can allocate maximum resources to push clients thus ensuring high fidelity.

Thus, by varying the allocation of resources, a server can either ensure high fidelity or greater scalability.

## V. Beyond PaP and PoP: PoPoPaP

PoPoPaP is a combination of the PoP and PaP approaches:

- The PaP alternative is added to a PoP server: To keep things simple, we simply replaced Push clients by PaP clients. Thus, (a) the average resiliency offered is better than PoP, (b) the degradation in service is also more graceful, and (c) the average coherency offered is higher than with PoP alone.
- Integration of PaP with PoP makes the approach more adaptive (by fine tuning using the parameters of PaP).
- The PoP algorithm at the server improves the average scalability offered by the system.

Since the resource requirements for Push as well as PaP are almost the same, the load at the server is the same as that of PoP. It should be noted that PaP generates pull requests (because server allows client to pull for $\epsilon$ time), which was absent in PoP. These pull requests are in addition to the requests generated by Pull clients and hence the load on the server may increase. But, as long as $\epsilon$ is small or the number of PaP clients is small as compared to Pull clients, this load is negligible and we get results similar to that of pure PoP.

As the proportion of pure-pull clients decreases, the extra load generated by replacing push with PaP also increases. For a very high ratio of PaP clients to pull clients, the number of extra pull requests generated by the PaP clients can be quite high. This leads to higher (pull) request denial rate, and fidelity can be compromised. Thus, the improvement in resiliency is achieved at the extra cost of a higher bandwidth utilization, higher server load and reduced fidelity. To understand the fidelity and overhead characteristics

TABLE III
Average Fidelity and Bandwidth Utilization of PoPoPaP

| **#Total Clients** | **#Pull Clients** | **PoP** | | | **PoPoPaP** | | | **Drop in Fidelity** | **Increase in Msg Overhead** |
|---|---|---|---|---|---|---|---|---|---|
| | | *Avg Fidelity* | *#Pulls* | *#Pushes* | *Avg Fidelity* | *#Pulls* | *#Pushes* | | |
| 224 | 176 | 87.50 | 15037 | 10462 | 83.03 | 21227 | 5434 | 4.47 | 18.14 |
| 224 | 152 | 90.11 | 12000 | 15010 | 83.15 | 20700 | 8500 | 6.96 | 27.91 |
| 224 | 128 | 92.03 | 9917 | 18270 | 84.50 | 19800 | 10987 | 7.53 | 32.76 |
| 224 | 104 | 93.40 | 7774 | 19717 | 84.60 | 18560 | 12050 | 8.80 | 39.43 |

of PoPoPaP, we conducted experiments by replacing the push component of a PoP server with PaP and compared them against PoP. The last two columns in table III summarize the results. The drop in fidelity

refers to the reduction in the average fidelity offered by PoPoPaP compared to the average fidelity offered by PoP. Similarly, the increase in message overhead indicates the percentage increase in the communication overhead (counting each pull as 2 messages and each push as 1 message) with PoPoPaP.

The first column corresponds to the total number of clients which is kept constant and the second column indicates the number of pure-pull clients.

PoP has better fidelity and lower message overheads compared to PoPoPaP as the proportion of Pure-Pull clients decreases (i.e., as we go down the table). These results are quite intuitive. In PoP, for a large number of push sockets, most of the clients enjoy push service. So the average fidelity offered and communication overheads are optimum (recall that push offers optimum QoS). When these push clients are shifted to PaP, the fidelity drops slightly. Also, because PaP also generates pull requests, communication overheads increase. These added pull requests load the server, as a result of which fidelity decreases. In summary, while PoPoPaP has the potential to offer higher resiliency, this comes at the cost of drop in fidelity and increase in the communication overheads.

Together, these arguments motivate the properties of PoPoPaP mentioned in Table I: by adaptively choosing pull or PaP for its clients, PoPoPaP has the ability to achieve the temporal coherency, scalability desired for a system.

## VI. Related Work

Several research efforts have investigated the design of push-based and pull-based dissemination protocols from the server to the proxy, on the one hand, and the proxy to the client, on the other. Push-based techniques that have been recently developed include broadcast disks [1], Salamander [18], continuous media streaming [3], publish/subscribe applications [19], web-based push caching [13], and speculative dissemination [4]. Research on pull-based techniques has spanned the areas of web proxy caching and collaborative applications [5], [6], [22]. Whereas each of these efforts has focused on a particular dissemination protocol, few have focused on supporting multiple dissemination protocols in web environment.

Netscape has recently added push and pull capabilities to its Navigator browser specifically for dynamic documents [20]. Netscape Navigator 1.1 gives two new open standards-based mechanisms for handling dynamic documents. The mechanisms are (a) *Server push*, where the server sends a chunk of data; the browser displays the data but leaves the connection open; whenever the server desires, it continues to send more data and the browser displays it, leaving the connection open; and (b) *Client pull* where the server sends a chunk of data, including a directive (in the HTTP response or the document header) that says "reload this data in 5 seconds", or "go load this other URL in 10 seconds". After the specified amount of time has elapsed, the client does what it was told – either reloading the current data or getting new data. In server push, a HTTP connection is held open for an indefinite period of time (until the server knows it is done

sending data to the client and sends a terminator, or until the client interrupts the connection). In client pull, HTTP connections are never held open; rather, the client is told when to open a new connection, and what data to fetch when it does so. Server push is accomplished by using a variant of the MIME message format "multipart/mixed", which lets a single message (or HTTP response) contain many data items. Client pull is accomplished by an HTTP response header (or equivalent HTML tag) that tells the client what to do after some specified time delay. The computation, state space and bandwidth requirements in case of Server push will be at least as much as was discussed in section II-C. In addition, since we are using HTTP MIME messages, the message overhead will be more (on average MIME messages are bigger than raw data). Because of the same reason, it is not feasible to use this scheme for highly dynamic data, where the changes are small and occur very rapidly. Also, it would be very difficult to funnel multiple connections into one connection as is possible in our model 1 (see equation 1). This will clearly increase the space and computation requirements at the server. For the Client pull case, for reasons discussed in Section VI-B, it is very difficult to use this approach for highly dynamic data. Still, these mechanisms may be useful for *implementing* the algorithms discussed in this paper as they are supported by standard browsers.

Turning to the caching of dynamic data, techniques discussed in [15] primarily use push-based invalidation and employ dependence graphs to track the dependence between cached objects to determine which invalidations to push to a proxy and when. In contrast, we have looked at the problem of disseminating individual time-varying objects from servers to proxies.

Several research groups and startup companies have designed adaptive techniques for web workloads [2], [5], [11]. Whereas these efforts focus on reacting to network loads and/or failures as well dynamic routing of requests to nearby proxies, our effort focuses on adapting the dissemination protocol to changing system conditions.

The design of coherency mechanisms for web workloads has also received significant attention recently. Proposed techniques include strong and weak consistency [16] and the leases approach [8], [23]. Our contribution in this area lie in the definition of temporal coherency in combination with the fidelity requirements of users.

Finally, work on scalable and available replicated servers [24] and distributed servers [7] are also related to our goals. Whereas [24] addresses the issue of adaptively varying the consistency requirement in replicated servers based on network load and application-specific requirements, we focus on adapting the dissemination protocol for time-varying data.

We end this section with a detailed comparison of two alternatives to our approaches: Leases [8], [10], a technique that also combines aspects from pull-based and push-based approaches, and Server-based prediction (instead of our client-based prediction) for setting Time-to-Live attributes for Web objects [9], [14], [16].

### A. Comparison with Leases

In the leases approach, the server agrees to push updates to a proxy so long as the lease is valid; the proxy must pull changes once the lease expires (or renew the lease). Thus, the technique employs push followed by pull. In contrast, the *PaP* approach simultaneously combines both push and pull—most changes are pulled by the proxy, changes undetected by the proxy are pushed to it. The leases approach has high fidelity so long as the lease is valid and then has the fidelity of pull until the lease is renewed. As shown earlier, by proper tuning, the fidelity of the *PaP* algorithm can approach that of push. The leases approach is more resilient to failures than a push (the duration of the lease bounds the duration for which the $tcr$ can be violated; the lease can renewed thereafter). The *PaP* approach has even greater resiliency than leases, since proxies continue to pull even if the server stops pushing. Finally, we note that the leases approach can be combined with the *PaP* algorithm—the lease duration then indicates the duration for which the server agrees to push "missed" (i.e., undetected) changes.

### B. Client Prediction vs. Server Prediction

PaP and PoP are based on using prediction capabilities at the clients/proxies. An alternative of course is to leave the prediction to the server. Such schemes are discussed in [9], [14], [16]. They use the `if-modified-since` field associated with the HTTP GET method (also known as conditional GET), together with the TTL (time-to-live) fields used in many of proxy caches. These schemes in general work as follows:

- Client does not use any TTR or prediction algorithm, but instead depends on some meta information associated with the data to decide the time at which to refresh the data.
- Since the server has access to all the data, it can use a prediction algorithm to predict a time when the data is going to change by $tcr$. The server then attaches this time value with outgoing data. Client will use this meta information to decide when to poll next. There is no need of a push connection.
- Since server has better access to data than client, server predictions will be in general more "accurate" than using a TTR algorithm at the client.

Though the *Server-Prediction* approach looks like a better option than PaP, it runs into following problems:

- the approach requires that previous history for the relevant data be maintained at the server. This will imply increased state information and computational needs at the server and will consequently adversely affect the scalability. Since in PoP (section IV) we reserve the pull method to serve clients when faced with problems of scalability, we prefer to make Pull relatively "light weight".
- the approach is more suitable for data that changes less frequently (e.g., say once every few hours). We are interested in web data that is highly dynamic and inherently unpredictable (e.g., data that changes every

few seconds/minutes such as stock quotes). For dynamic data, the performance will be slightly better than Adaptive TTR, but at a cost of server resources and scalability.

- if the server prediction is wrong and still a change of interest occurs in data, the server is helpless since it cannot push the change to the client. The change is lost. This will never happen in PaP.

In summary, so far dynamic data has been handled at the server end [15], [9]; our approaches are motivated by the goal of offloading this work to proxies.

## VII. Concluding Remarks

Since the frequency of changes of time-varying web data can itself vary over time (as hot objects become cold and vice versa), in this paper, we argued that it is *a priori* difficult to determine whether a push- or pull-based approach should be employed for a particular data item. To address this limitation, we proposed two techniques that combine push- and pull-based approaches and adaptively determine which approach is best suited at a particular instant.

- Our first technique (*PaP*) is inherently pull-based and maintains soft state at the server. The proxy is primarily responsible for pulling those changes that are of interest; the server, by virtue of its soft state, may optionally push additional updates the proxy, especially when there is a sudden surge in the rate of change that is yet to be detected by the proxy. Since the server maintains only soft state, it is neither required to push such updates to the proxy, nor does it need to recover this state in case of a failure.
- Our second technique (*PoP*) allows a server to adaptively choose between a push- or pull-based approach on a per-connection basis (depending on observed rate of change of the data item or the coherency requirements).

We also showed how PoP can be extended to use PaP for some of its connections.

A key contribution of our work is the design of algorithms that allow a proxy or a server to efficiently determine *when* to switch from a pull-based approach to push and vice versa. These decisions are made based on (i) a client's temporal coherency requirements ($tcr$), (ii) characteristics of the data item, and (iii) capabilities of servers and proxies (e.g., a pure HTTP-based server precludes the use of push-based dissemination and necessitates the use of a pull-based approach by a proxy).

Our techniques have several characteristics that are desirable for time-varying data: they are *user-cognizant* (i.e., aware of user and application requirements), *intelligent* (i.e., have the ability to dynamically choose the most efficient set of mechanisms to service each application), and *adaptive* (i.e., have the ability to adapt a particular mechanism to changing network and workload characteristics). Our experimental results demonstrated that such tailored data dissemination is essential to meet diverse temporal coherency requirements, to be resilient to failures, and for the efficient and scalable utilization of server and network resources.

Currently we are extending the algorithms developed in this paper to design algorithms suitable for cooperative proxies and also for disseminating the results of continual queries [17] posed over dynamic data.

## REFERENCES


[1] S. Acharya, M. J. Franklin and S. B. Zdonik, Balancing Push and Pull for Data Broadcast, *Proceedings of the ACM SIGMOD Conference, May 1997.*

[2] FreeFlow Product Details, Akamai Inc., *http://www.akamai.com/service/freeflow.html, 1999.*

[3] E. Amir, S. McCanne and R. Katz, An Active Service Framework and its Application Real-time Multimedia Transcoding., *Proceedings of the ACM SIGCOM Conference, pages 178-189, September 1998.*

[4] A. Bestavros, Speculative Data Dissemination and Service to Reduce Server Load, Network Traffic and Service Time in Distributed Information Systems., *International Conference on Data Engineering, March 1996.*

[5] P. Cao and S. Irani, Cost-Aware WWW Proxy Caching Algorithms., *Proceedings of the USENIX Symposium on Internet Technologies and Systems, December 1997.*

[6] A. Chankhunthod, P. B. Danzig, C. Neerdaels, M. F. Schwartz and K. J. Worrel A Hierarchical Internet Object Cache., *Proceedings of the 1996 USENIX Technical Conference, January 1996.*

[7] D. M. Dias, W. Kish, R. Mukherjee and R. Tewari A Scalable and Highly Available Server., *Proceedings of the IEEE Computer Conference (COMPCON), March 1996.*

[8] V. Duvvuri, P. Shenoy and R. Tewari, Adaptive Leases: A Strong Consistency Mechanism for the World Wide Web. *InfoCom* March 2000.

[9] V. Cate Alex - A Global File System. *Proceedings of the 1992 USENIX File System Workshop*, pages 1-12, May 1992.

[10] C. Gray and D. Cheriton, Leases: An Efficient Fault-Tolerant Mechanism for Distributed File Cache Consistency, *Proceedings of the Twelfth ACM Symposium on Operating System Principles, pages 202-210, 1989.*

[11] A. Fox, Y. Chawate, S. D. Gribble and E. A. Brewer, Adapting to Network and Client Variations Using Active Proxies: Lessons and Perspectives., *IEEE Personal Communications, August 1998.*

[12] S. Gribble, E. Brewer, System Design Issues for Internet Middleware Services: Deduction from Large Client Trace., *Proceedings of the USENIX Symposium on Internet Technologies and Systems, December 1997.*

[13] J. Gwertzman, M. Seltzer, The Case for Geographical Push Caching., *Proceedings of the 5th Annual Workshop on Hot Operating Systems, pages 51-55, May 1995.*

[14] J. Gwertzman, M. Seltzer, World-wide Web Cache Consistency., *Proceedings of 1996 USENIX Technical Conference, January 1996.*

[15] A. Iyengar and J. Challenger, Improving Web Server Performance by Caching Dynamic Data., *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USEITS), December 1997.*

[16] C. Liu and P. Cao, Maintaining Strong Cache Consistency in the World-Wide-Web., *Proceedings of the Seventeenth International Conference on Distributed Computing Systems, pages 12-21, May 1997.*

[17] L. Liu, C. Pu and W. Tang, Continual Queries for Internet Scale Event-Driven Information Delivery., *IEEE Trans. on Knowledge and Data Engg., July/August 1999.*

[18] G. R. Malan, F. Jahanian and S. Subramanian, Salamander: A Push Based Distribution Substrate for Internet Applications., *Proceedings of the USENIX Symposium on Internet Technologies and Systems, December 1997.*

[19] A. G. Mathur, R. W. Hall, F. Jahanian, A. Prakash and C. Rasmussen, The Publish/Subscribe Paradigm for Scalable Group Collaboration Systems., *Technical Report CSE-TR-270-95, Dept. of Computer Science and Engg., University of Michigan, 1995.*

[20] An Exploration of Dynamic Documents, Netscape Inc., *http://home.netscape.com/assist/net_sites/pushpull.html.*

[21] V. Paxon and S. Floyd, Wide Area Traffic: The Failure of Poisson Modeling., *IEEE/ACM Transactions on Networking, pages 226-244, June 1995.*

[22] Raghav Srinivasan, Chao Liang and Krithi Ramamritham, Maintaining Temporal Coherency of Virtual Data Warehouses, *The 19th IEEE Real-Time Systems Symposium (RTSS98), Madrid, Spain, December 2-4, 1998.*

[23] J. Yin, L. Alvisi, M. Dahlin and C. Lin, Hierarchical Cache consistency in a WAN., *Proceedings of the USENIX Symposium on Internet Technologies and Systems, October 1999.*

[24] H. Yu and A. Vahdat, Design and Evaluation of a Continuous Consistency Model for Replicated Services., *Proceedings of OSDI'2000, October 2000 (to appear).*