Accelerating Dynamic Web Content Delivery Using Keyword-based Fragment Detection

Daniel Brodie, Amrish Gupta, and Weisong Shi

Department of Computer Science
Wayne State University
Detroit, Michigan 48202

Abstract. The recent trend in the Internet traffic is increasing in requests for dynamic and personalized content. To efficiently serve this trend, several server-side and cache-side fragment-based techniques, which exploit reuse of Web pages at the sub-document level, have been proposed. Most of these techniques do not focus on the creation of the fragmented content from existing dynamic content. Also, existing caching techniques do not support fragment movement across the document, a common behavior in dynamic content.

This paper presents two proposals that we have suggested to solve these problems. The first, DyCA, a dynamic content adapter, takes original dynamic Web content and converts it to fragment-enabled content. Thus the dynamic parts of the document are separated into separate fragments from the static template of the document. This is dependent on our proposed keyword-based fragment detection approach that uses predefined keywords to find these fragments and to split them out of the core document. Our second proposal, an augmentation to the ESI standard, allows splitting the information of the position of each fragment in the template from the template data itself by using a mapping table. Using this, a fragment enabled cache can have a more fine grained level of identifying fragments independent of their location on the template, which enables it to take into account fragment behaviors such as fragment movement.

We used the content taken from three real Web sites to achieve a detailed performance evaluation of our proposals. Our results show that our keyword-based approach for fragment extraction provides us with cacheable fragments that, when combined with our proposed mapping table augmentation, can provide significant advantages for fragment-based Web caching of existing dynamic content.

1 Introduction

Researchers have recently proposed several server-side and cache-side mechanisms to improve the generation and serving of dynamic Web content. Server-side techniques, exemplified by techniques such as delta encoding [1], data update propagation [2], fragment-based page generation [3, 4], reduce the load on the server by allowing reuse of previously generated content to serve new requests. Cache-side techniques, exemplified by systems such as Active Cache [5], Gemini [6], CONCA [7], and the content assembly technique proposed by Wills *et al.* [8], attempt to reduce the latency of dynamic content delivery by moving some functionality to the edge of network. Similar trends are also visible in commercial caching and edge server products, most notably

IBM's WebSphere [9] and Akamai's Edgesuite [12]. Despite their difference in focus, both server-side and cache-side approaches share the same rationale, specifically that it is possible to view the document in terms of a quasi-static *template* (expressed using formatting languages such as XSL-FO [10] or, what is currently becoming the *de facto* standard, edge-side include (ESI) [11]), which is filled out with multiple individually cacheable and/or uncacheable *objects*[1]. This object composition assumption enables surrogates and downstream proxy caches to reuse templates and cached objects to efficiently serve subsequent requests and additionally reduce server load, bandwidth requirements, and user-perceived latencies by allowing only the modified or unavailable objects to be fetched.

Although the above techniques appear promising, there are a number of issues that are not addressed in these current infrastructures. Even though there might techniques used by certain companies [12], due to their closed nature we cannot check them, and so to the best of our knowledge, there is no open and free method of separating objects from existing dynamic document, except from our existing work on DYCE [13]. Also, current technologies for supporting dynamic objects do not differentiate between the location of the objects in the document, and object itself. This makes it impossible to efficiently implement the situations where the object can move between different places in the document without changing data, which is common in certain news Web sites [17].

This paper describes our efforts on addressing these shortcomings. We are proposing two methods that should solve these shortcomings. The first is an augmentation to the ESI standard, the most used method for specifying the format of the templates, to allow the fragment locations to be specified in a mapping table that is sent with the template. This allows the objects to move across the document without needing to re-serve the template. Our second proposal, DyCA, a Dynamic Content Adapter, is a two part model for creating object-based content from original dynamic content. The first part extracts the objects from the original content, giving us the needed separation between template, objects, and object location by using our mapping table approach. The second part of DyCA delivers the content to a fragment-enabled client, like a caching proxy server. A Python-based fragment-enabled caching proxy, named CONCA-Lite was developed to allow the testing of the object extraction, and content delivery modules of DyCA.

Our method for creating the dynamic content from the original Web content is based on a simple and effective keyword based object extraction technique to find dynamic objects inside a static Web page. The dynamic content can then be served by our DyCA server, which can serve fragments from the document as needed, enabling a client to support template and object caching. Our proposed ESI-extended format allows for caching of both the objects and the template and allows for object movement. This type of concept, where the object in the template maps, based on a mapping table, is, to our knowledge, introduced here for the first time. By having a fully functional fragment-enabled content server and client, and by testing on real world data, we have gotten accurate results, beyond regular experiments done in the field. These results have shown that our proposed method for fragment extraction based on the keywords in the

---

[1] The terms *objects* and *fragments* will used interchangeably in this paper.

document can allow us to cache existing non-fragmented content and achieve significant performance improvements by utilizing our proposed augmentation for a mapping-table based template.

The rest of the paper is organized as follows. Section 2 describes the related background for this field. Section 3 addresses problems with the current infrastructures. Section 4 shows the design and implementation of the system. Section 5 presents the evaluation and results of the testing of our architecture. Section 6 concludes the paper and discusses our planned future work.

2 Background

2.1 Fragment Based Caching
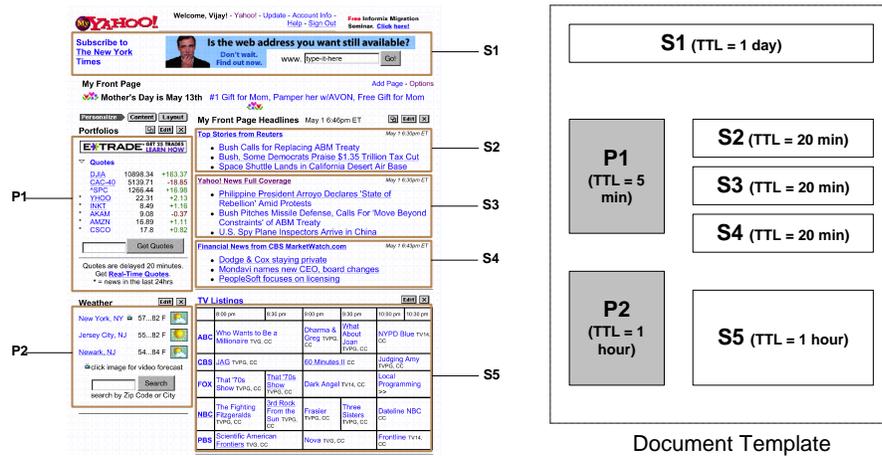
Fig. 1. Dynamic content can be viewed in terms of a quasi-static document template and individual objects, which exhibit different sharing, cacheability, and freshness time characteristics.

One fundamental block in caching dynamic content is allowing to split up a document into different static and dynamic parts. By doing this, parts of the document, called *fragments*, can be treated separately rather then treating the document as a whole. Thus each fragment can have its own behavior allowing more fine-grained control over caching behavior and data sharing of the different segments of a document. Certain fragments of a document can be shared between different clients whereas some clients want different information in other fragments. Also, some parts of a document change more frequently than others, while other parts are completely static. By treating the document as a whole rather then separating it into fragments a page cannot be partially shared between users nor can it be partially cached. Rather, when a little part of the document changes the whole document needs re-fetching, and if parts of the document can't be shared, then the document can't be shared at all. For example, consider a popular customizable Web site with content that gets continuously updated with information such as news and weather. Figure 1 shows the snapshot of a personalized my.yahoo.com

page, which fits in such an example, and what the corresponding document template and component objects might look like. `S*` and `P*` represent objects that are shared and private respectively, and `TTL` captures the length of time this object remains valid. The contents of fragments on the page can change as new stories develop or as the weather changes, leaving other fragments unchanged and with no need to re-fetch the data. Users that are viewing sports stories can share those fragments, and if some of those users are also viewing the economic section, then that can be shared by all the people viewing that fragment as well. There currently are a number of ways to split a document into fragments and reconstruct it. Different approaches do reconstruction in different parts of the network, including the originating server, the cache proxy, and the client. Each one of these methods has a different way of dealing with caching and the fragments.

2.2 Edge-Side Includes

The ESI [11], Edge-Side Include, has currently become the de facto standard in specifying the format for templates in fragment based documents. It uses a simple XML-based markup format that specifies content fragments for inclusion and dynamic assembly in a base template. These ESI specific tags are provided as extensions to the traditional HTML format allowing for minimal format change to the original document format. It separates the document in such a way that allows for the server or proxy to manage the objects as separate entities. This allows for different levels of cacheability for each fragment, and for large amount of dynamic content to be cached. ESI also includes a very complete framework for conditional getting of fragments, cookie support, and error control. Every `esi:include` tag references a specific URI with a TTL, all of this is included in the template file which gives the layout and aesthetic look to the document. Thus, all the information regarding the fragments and the template are actually sent in the template itself. For a client to support the ESI framework all that it needs to do is to implement support for parsing and acting based on the template format. Thus, ESI's simplicity is a very strong point.

CONCA [7] allows for additional client information, such as the type of device of the client, to modify the returned content, based on a different template yet reusing the same objects.

3 Existing Problems of ESI

An important factor of the efficiency of fragment based documents is the method used to update it. One of the popular uses in fragment documents is for object movement. This type of behavior, as can be seen in Figure 2, is represented when one or more fragments from a dynamic document move between the different available positions on the template. A good example is a news Web site where old stories (represented as fragments) move down the document and new ones are added from the top. If we use ESI to construct our document then there are two possible ways to update the document so that the objects can move across the document. Since in ESI the document is split up into two parts, the template, and the fragments, then one possibility will be to update the fragments where the object moved to where it moved from. Another possibility is
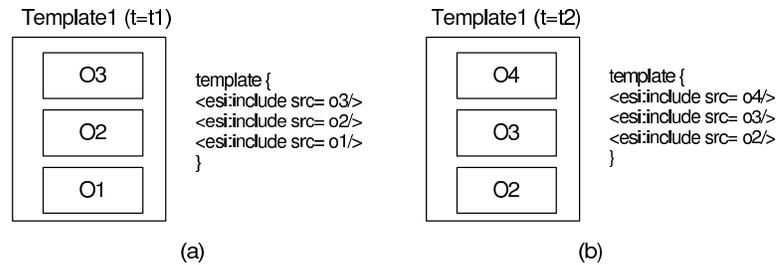
Fig. 2. An example of template caching problem.

to update the template such that the updated template has the URL of the new locations in the template pointing to the correct fragment.

3.1 ESI - Static Template

One obvious solution to fragment movement, which we will call *static template*, is to update the fragments themselves, leaving the template completely static. Using our news Web site example from earlier, the news template points to a series of objects, and when a new story gets added to the page it pushes the old one out of the page. All the objects in the page need to expire, be invalidated, and be re-retrieved as the fragments lower down on the page. Considering that a document, like a news Web site, that has most of its objects moving around the Web page, this means invalidating most of the objects on the page and fetching them from the server on a regular basis. The biggest problem with this method is that most of these fragments are already present on the client just in a different location, and re-fetching them in a different place means a lot of wasted data transfer that could otherwise be cached.

3.2 ESI - Static Objects

A different approach to solve this problem, which we will call *static objects*, is to in-validate the template and re-fetch it for spatial changes, leaving the objects static for such a change. The fragments will still contain only dynamic data and would need to be re-fetched for a data change. The new template will have the object in the first position pointing to the new objects, and all the other positions pointing to the moved objects. This seems to solve the problem since the objects are unmodified do not need to be re-fetched. Supposedly, this saves a lot of data transfer by transferring only the new objects and not requiring all the objects to be re-fetched as in our earlier example. Since the only things that has actually changed in the template is the URL of a few of the fragments, this means that most of the transferred template, which would still contain only static data, is already on the client. Considering that the template can be very large, as we show later in section 5, doing this on a regular basis, and transferring all this data that could otherwise be cached, this is not a very efficient solution as it might seem at first.

3.3 Mapping Table

There is no way to solve the problem of object movement in the current ESI infras-tructure. Either you will be invalidating many objects that really are valid, or you will

be invalidating a template which has very little data actually modified. That is why a proposed solution needs to make additions to the ESI infrastructure. These additions, a mapping table that gets sent along with the template and an addition to the template format to allow inclusion of objects from the mapping table, allow for a new method of fragment-based caching, which we will denote as *mapping table*, which does not requires either the objects or the templates to be invalidated for spatial changes. Thus, when a client retrieves a template, the client will also cache the mapping table that was sent with the template. When this client needs to fetch a fragment from the template, the identifier in the template is looked up in the mapping table, and then the fragment is fetched from the appropriate URL. The mapping table is relatively a small amount of data compared to the template and object sizes. When an object needs to be moved across a document, from one locations in the template into another, the only thing that needs to be updated is the mapping table. Thus, both the template and the objects continue to be cached and treated efficiently, while still allowing for object movement. The mapping table's small size and ease of use makes it optimal for these cases.
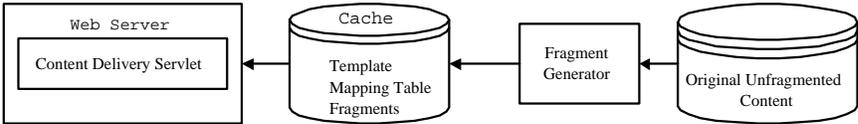
4   DyCA: Dynamic Content Adapter

Fig. 3. The general architecture of DyCA.

4.1   Design

As mentioned earlier, DyCA, our proposed **Dy**namic **C**ontent **A**dapter, has been designed to augment the existing servers to serve dynamic content efficiently. In the client-proxy-server model it sits on the front-end of the server and can then take existing Web content that is requested from the server and process it to serve the dynamic content instead. So when the original content changes, DyCA will regenerate its fragment-enabled content. This allows it to be deployed anywhere from the same location as the actual Web server or to the ISP level. As shown in Figure 3, DyCA is actually split up into 2 separate but very important parts, the Dynamic Content Generator module, and the Dynamic Content Delivery module. The generator module deals with taking the existing dynamic content and converting it into fragmented content. The delivery module can then take the dynamic Web pages generated by the earlier module, and serve them to the client appropriately.x These two modules together let us take an existing static Web site and easily change it into a fragment-based dynamic Web site, that can be cached properly.

The content generation module generates the objects and template from the original dynamic Web content. It uses a keyword based approach to split up the dynamic content into fragments. The keyword based approach works by building an XML of the Web site and searching this XML tree for specific tags that can signify a different object. These

tags are set separately for different Web site based on the structure of the HTML and the content. By looking at popular Web sites, such as the personalized `my.yahoo.com` page shown earlier in figure 1, it is fairly simple to see the implicit fragments contained in the document. Fragments can be easily distinguished based on certain differences such as a different font or a table tag, and based on certain predefined keywords, such as the TV Listings headlines, or the Weather headlines. Once the tags in the XML tree are identified, the children XML tags and the rest of the XML content contained in the identified tags is extracted to create the objects. A special include tag, that has a special object id for each object, will be placed in the position where the object was extracted from the main document. The mapping table is just a list of the object ids and their corresponding URLs in a parse-able file.

Additional information for each object, such as the TTL of the object, is calculated by looking at a long term overview of the Web site. Numerous instances of the Web page are collected over a regular period of time. Each instance is parsed and separated into the different objects, template, and mapping table by the content generation module. By comparing the objects and their position over the period of time, an accurate dynamic behavior can be seen that allows the correct generation of the mapping table to be most efficient with regards to object movement. It also allows the TTL for each fragment location in a document to be generated. Since the location in the TTL of a location in the template remains mostly the same, this TTL is then reused later as the TTL for each fragment location in the document. In our news Web site example, the sidebar listing all the news from yesterday will always have a TTL of 24 hours.

The content delivery module is responsible for serving the template, objects, and mapping table to the fragment-enabled proxy cache. The content delivery module uses the data created by the content generation module. The content delivery module needs to implement the extensions to the protocols in order to send the data created by the content generation module in an appropriate way. It needs to add information regarding the mapping table, and so notify its client, the fragment-enabled proxy cache, when a request is dynamic and has a template or when it is static. By sending the mapping table to the client, the client can then get the static URLs of the objects to be able to access them from the template. The content delivery module also needs to support the client updating only its mapping table, so that bandwidth would not need to be wasted with already cached objects, or templates. This module can be backwards compatible with existing technologies and can support serving to both client level caching [14], and proxy level caching [7, 12].

4.2 Implementation

The dynamic content generator is program that parses existing Web sites and outputs a dynamic, fragment based, cacheable, Web page. Three Web sites, New York Times [15], India Times [16], and Slashdot.com [17], were chosen due to their dynamic nature, and since none of them supported any form of cacheable dynamic data. The Web sites were constantly monitored for changes during an extensive period covering 2 weeks. This data was then passed through the keyword based object extraction. Each object was extracted from each instance of the Web site by finding appropriate, predefined tags in the document. Once each object was extracted and the ESI-based template was constructed,

the resulting fragmented documents were compared across their time element to calculate the TTL and the mapping table. Object movement across the document is taken into account and allows for the object to not be replaced too soon, and remain in the mapping table. An object was considered expired once it wasn't in any part of the document. Once the template, objects, and mapping table are in place from the generator
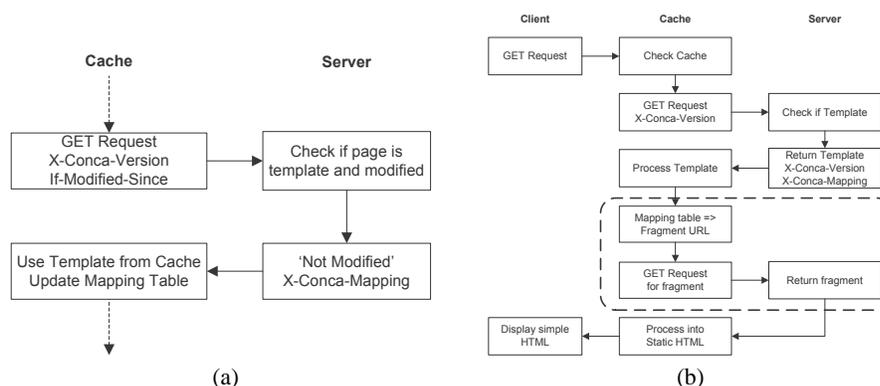
Fig. 4. (a) Process of initial retrieval (b) Process of cached retrieval.

module, the dynamic content delivery module just needs to serve them. Implemented as a server extension using servlets in Java and sitting on top of popular Web servers such as the Apache Web Server [18] or the Jigsaw Web Server [19], this module appears as a traditional Web server to regular clients, but provides the dynamic content ability to able clients. This module uses information from the generator module to build up information such as which pages are templates, the TTL for certain objects, and the mapping table. Figure 4 shows the process of events when a client request for a Web page. When a CONCA-enabled request from a proxy for a document is made, if the document has fragments then the mapping table is looked up and added as an `X-CONCA-MAPPING` HTTP header to the response. The expiration of the object is set based on the TTL contained with the mapping table. The body of the response is just the ESI augmented template. Once the cache has the template, it will go on and request the fragment as needed from the server. The cache can then build up the proper final document and send that off to the client. The servlet's support of the `If-Modified-Since` HTTP header when requesting the template is crucial to the efficiency of the mapping table. When the template expires on the cache, the cache will request the template over again using the `If-Modified-Since` HTTP header. Since the template rarely changes, this will mean that the cache will usually get a '304 Not Modified' response. This response will contain the new mapping table, allowing the cache to update it's mapping table without having to re-transfer the template. If a static document, or a static fragment is requested from the content delivery module, then the content delivery module will behave just like a regular Web server.

4.3 ESI Augmentation

As explained earlier, ESI provides and extensive range of existing technology to support a wide range of uses in structuring a template file and specifying things like TTL and so forth. In trying to remain as standard compliant as possible, the ESI format was picked to represent the structure of the template. Yet the ESI standard only supports document fragments identified by static URLs, which will not suffice in our case. Thus we needed to augment the ESI standard by adding the `esi:xconca-include` tag. This tag allows the specification of an ID number that can be looked up by the client in the mapping table and retrieve the object's static URL.

5 Performance Evaluation

The experimentation of our proposed method for object extraction and object delivery required simulating a fragment-enabled client-cache-server model. Using this model we can compare different types of performance for the different types of fragment-enabled dynamic content behaviors. The experiments we used to test the performance of the different caching systems targeted user-perceived latency and bandwidth usage specifically.

5.1 CONCA-Lite

The experiments in our client-cache-server model required a proxy that supported our proposed augmentation to the ESI and supported the dynamic assembly of the final content for the client. To achieve this, a simple cache proxy, called CONCA-Lite, was implemented using Python and its `asyncore` modules to create a simple extensible proxy. It was designed such that testing different caching methods would require little or no change on the proxy side. Thus, allowing us to make fair and accurate comparisons between the different caching methods, which are tested on the same caching framework. This CONCA-Lite proxy, which implemented a minimalistic version of the CONCA proposal [7], was then used to test the different dynamic caching methods described later.

5.2 Evaluation Platform

We simulated the client-cache-server model using three machines, all connected on a 100Mbit/s LAN, at most, separated by a switch. The server, a 2.0 Ghz Pentium 4 machine with 512 MB RAM with Linux, ran the Jigsaw Java server to host the DyCA servlets. The cache, a 2.4 Ghz Pentium 4 machine with 1 GB RAM with Linux, ran the CONCA-Lite. The client, a 2.2 Ghz machine with 512B of RAM with Linux, ran a Python-based simulation of a client accessing Web pages in a predefined order.

We modeled 4 different caching behaviors using our experimentation. To test each approach, the client was set up to request the Web page of the server from the cache at request intervals of 10 seconds, for a total of 10 minutes. The cache would check to see if it has the needed document, request the document from the server if it needs to, and return the document to the client. When testing a caching behavior that has a fragment enabled template, it would request all of the objects in the template, it would construct

the final document, and return it to the client. Thus the client does not need to implement anything beyond the standardized HTTP protocol. The first method, using no fragment caching, was implemented by disabling caching in the proxy, and having the server send the original Web page. This is consistent with the behavior of real dynamic content using static pages in today's Internet, due to cookies, and other such information, that render a page uncacheable. In the second method the template of the document remains static, while the fragments of the objects are updated for content change. The template was cacheable for the whole testing session, while the objects were cacheable for as long as their TTL was valid. In the third caching behavior, the template is updated when a fragment moves between locations in the document, and the objects change due to data changes only, and not spatial changes. The last model represents our proposed mapping table approach. When the template is returned a mapping table is returned with it in the HTTP header, the proxy can then cache the mapping table, and update the mapping table when a spatial change happens. The template remained static for the testing session, while the objects only changed for data changes.
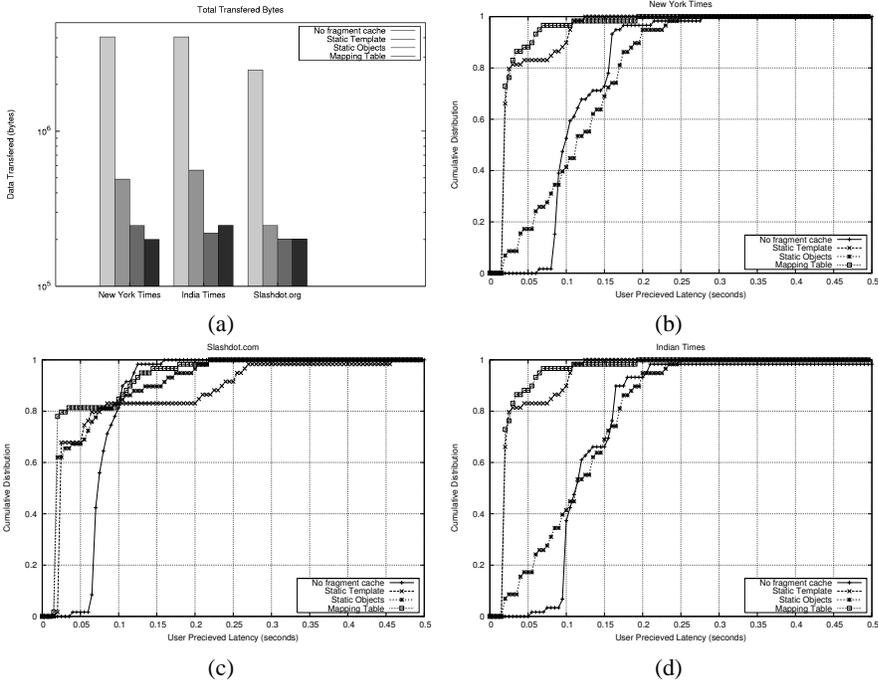
Fig. 5. Evaluation results: (a) total data transfer between server and cache, and user perceived latency for: (b) New York Times, (c) Slashdot, and (d) India Times.

5.3 Experimentation

Three Web sites, New York Times [15], India Times [16], and Slashdot.com [17], were used for testing each approach. Two types of measurements were taken during the testing to evaluate the performance, the total amount of data transfered between the client

and the server, and the user perceived latency per user request. The first type of measurement is important to show what method performs best as a cache, with the least data transfer between the cache and the server. There is no need to check for the data transfer between the client and the cache, since in all four models it should remain roughly the same. The generation of the final non-fragmented Web page by the cache that gets sent to the client, and the method it is updated, is what changes between each method. The second measurement type is important to show how an improved Web caching architecture will benefit the client as well, and not only the server.

5.4 Results and Analysis

Figure 5(a) demonstrates the total bytes used in transferring data in all of the methods mentioned. As can be seen from the graphs, using a static template and updating the objects to support object movements requires considerable more data transfer between the server and the proxy. This is the method that is most commonly used today in dynamic Web sites. Using static objects and a dynamic template to achieve a dynamic Web page might seem efficient enough when looking at the bytes transfered, but, as we will see later, this efficiency is lost when looking at the user perceived latency. There is no comparison, though, between any of the dynamic methods and using traditional static objects. The amount of data transfer when not using a fragment based architecture can be more then 10 times the amount of data transfer when using a good fragment based caching. When using a mapping table to transfer the data, the amount of bytes transfered is considerably smaller. In fact, the total bytes transfered with a mapping table is little over the total size of the initial site and the size of the changes, meaning a minimal amount of wasted data is being transfered. This should considerably reduce the server load when using such an architecture. Figures 5 (b)(c)(d), show the user perceived latency in seconds for New York Times, Slashdot, and India Times respectively. This figure shows that the user has to wait the least amount of time for the Web site when the mapping table architecture is used. With regard to user perceived latency, the static object method's performance is almost as bad as the performance of using regular static Web pages. The only method that comes close to the method of using a mapping table is the static template method, which as we saw before performed badly when looking at the amount of bytes transfered. This type of optimization is very important for the client so it may receive its data in a timely manner, especially clients that use slower connections such as dialup. Otherwise, from the users prospective, the actual retrieval of the page is slower. From these figures we can conclude that the mapping table method has performed better then all of the other proposed methods in all of our tested fields.

We can see some unaccounted behavior in how the 2nd and 3rd method flip in their efficiency between the results of the amount of bandwidth used, and the user latency done. When the template is static and there is no mapping table, the objects get transfered at a higher request rate since the proxy can't cache them due to object movement. This extra data transfer has little effect on the user perceived latency in our testing environment due to it being a high speed network. Yet this extra data transfer is significant in terms of amount of data transfered, as can be seen in the earlier figure. Since the only thing that needs to get updated every once in a while is the transferring of the template, in terms of latency, this is very close to transferring a mapping table at about the same

interval. Yet by looking at Table 1 you can see that the template is considerably larger then the mapping table in most cases. This is what causes the large amount of bytes to be transfered. Had we artificially slowed down the network, the user perceived latency for static objects would have been much greater. In the case of static objects, the templates is considered dynamic, and gets updated every time there an object moves. Every such template fetch requires the cache to re-parse the template, which can be very large as seen in Table 1, and recheck it's cache for every single object, in some cases this requires the cache to send HTTP request to see if the data was modified. This type of overhead causes the extra latency seen in the graph.

	New York Times	India Times	Slashdot
Template Size	17 KB	15 KB	1.7 KB
Avg. Object Size	3.6 KB	4.8 KB	0.6 KB
Mapping Table Size	1.0 KB	0.8 KB	2.2 KB

Table 1. The comparison of template and object sizes for the different Web sites.

6 Related Work

Dynamic Web content delivery have increasingly becoming an important Web engineering issue as more and more Web content are generated in a dynamic and personalized way. Fragment-based techniques have received considerable attention from the research community in recent years [2–4, 8, 14]. Most of these approaches either assume the fragment-based content is served by Web server automatically, or look at server-side caching only.

To our knowledge, few of existing work discuss the manner of how to generate fragment from existing legacy Web servers without server-side information. One of the first effort is DYCE [13], which is model-based dynamic Web content emulator. Recently, Ramaswamy *et al.* proposed a novel scheme to automatically detect and flag fragments [20], which share the similar goal of this paper. However, there are three differences between us: First, although both of our work intends to automatic detection of fragments, our keyword-based is simple and easy to implement, while their approach is complex and has theoretical analysis. However, which one is better is still not clear. Second, in our work we focus on engineering implementation of DYCA, while their work focuses on automatic detection. In this sense, their work is a good complement to DYCA. Third, the mapping table based fragment delivery proposed in this paper is novel.

Edge Side Includes [11] is becoming one of the foundation blocks in specifying a common format and method for fragments and templates in this field. It is popular among many different existing methods. Naaman *et al.* [21] have done studies comparing ESI to delta encoding, finding ESI to have possible performance advantages.

Automatic detection of templates from Web pages has been studied from data mining field as well [22, 23]. They discuss the problem of template detection through discovery of pagelets in the Web pages. However, our work differs from the work on template detection both in context and content. First, the work on template detection is aimed towards improving the precision of search algorithms. While our aim is improving dynamic content delivery. Second, only template is interested in their work,

while we care both template and fragments. Therefore, the method used in these two approaches are different too. The method presented there to finding fragments is done based on amount of hyperlinks present in certain parts of the document. They do not build up an XML tree, nor treat anything more then hyperlinks, unlike we have done in our approach. This method applies better to search algorithms rather then to dynamic fragment extraction.

Our current research differentiates from earlier work done on DYCE [13], the Dynamic Web Content Emulator. With DYCE, we were attempting to build up general models for usage to describe the behavior of current fragment-based caching. Although it looked promising at the time, it generated too many objects and didn't match up to actual real world designs. Our current research was an attempt to try and continue that same research using real world Web sites so as to get more correct results.

Other research groups [24, 25] have also defined other criteria for finding objects in documents. While they have focused on content of the fragments and of the Web pages themselves, we have focused on their existence on a spatial, and location axis in the document.

7 Conclusions and Future Work

We have shown our proposed solution for keyword based object extraction, and object delivery. We have also explained our proposal of augmenting the ESI to include support for a mapping table. We have implemented these proposals into DyCA and then by taking actual Web pages and running through DyCA's keyword-based extraction to transform them into fragment-enabled content we have been able to run simulations between our sample proxy and the DyCA adapter. These simulations allowed us to compares our proposals to the current available methods of serving dynamic content on the Web. Our keyword-based approach allows for creation of dynamic content in such a way as to maximize the cache-ability of the content in a fragment-enabled caching system. Using the mapping table approach in the cache proxy which, according to our results, will give the best performance for both the server and the client, together with our DyCA adapter we can effectively cache in an efficient way Web sites that currently use non-fragmented content. Currently our DyCA and CONCA-Lite implementations are very young, and could still be further optimized. Our future work consists of continuing testing of these implementations to further refine the design of our CONCA prototype [7], which incorporates a novel design for efficient caching of dynamic and personalized content.

References

1. Mogul, J.C., Douglis, F., a. Feldmann, Krishnamurthy, B.: Potential Benefits of Delta-Encoding and Data Compression for HTTP. In: Proc. of the 13th ACM SIGCOMM'97. (1997) 181–194
2. Challenger, J., Iyengar, A., Dantzig, P.: A scalable system for consistently caching dynamic Web data. In: Proc. of IEEE Conference on Computer Communications (INFOCOM'99). (1999)
3. Challenger, J., Iyengar, A., Witting, K., Ferstat, C., Reed, P.: A publishing system for efficiently creating dynamic Web content. In: Proc. of IEEE Conference on Computer Communications (INFOCOM'00). (2000)

4. Douglis, F., Haro, A., Rabinovich, M.: HPP:HTML macro-pre-processing to support dynamic document caching. In: Proc. of the 1st USENIX Symposium on Internet Technologies and Systems (USITS'97). (1997) 83–94

5. Cao, P., Zhang, J., Beach, K.: Active cache: Caching dynamic contents on the Web. In: Proc. of IFIP Int'l Conf. Dist. Sys. Platforms and Open Dist. Processing. (1998) 373–388

6. Myers, A., Chuang, J., Hengartner, U., Xie, Y., Zhang, W., Zhang, H.: A secure and publisher-centric Web caching infrastructure. In: Proc. of IEEE Conference on Computer Communications (INFOCOM'01). (2001)

7. Shi, W., Karamcheti, V.: CONCA: An architecture for consistent nomadic content access. In: Workshop on Cache, Coherence, and Consistency(WC3'01). (2001)

8. Wills, C.E., Mikhailov, M.: Studying the impact of more complete server information on Web caching. In: Proc. of the 5th International Workshop on Web Caching and Content Distribution (WCW'00). (2000)

9. IBM Corp.: http://www.ibm.com/Websphere (Websphere platform)

10. http://www.w3.org/Style/XSL/ (W3C XSL Working Group)

11. Tsimelzon, M., Weihl, B., Jacobs, L.: ESI language specification 1.0 (2000)

12. http://www.akamai.com/ (Akamai Technologies Inc.)

13. Shi, W., Collins, E., Karamcheti, V.: DYCE: A synthetic dynamic Web content emulator. In: Poster Proc. of 11th International World Wide Web Conference. (2002)

14. Rabinovich, M., Xiao, Z., Douglis, F., Kamanek, C.: Moving edge side includes to the real edge – the clients. In: Proc. of the 4th USENIX Symposium on Internet Technologies and Systems (USITS'03). (2003)

15. http://www.nytimes.com (Nytimes Web site)

16. http://www.indiatimes.com (Indiatimes Web site)

17. http://www.slashdot.com (Slashdot Web site)

18. http://httpd.apache.org (Apache HTTP Server Project)

19. http://www.w3.org/Jigsaw (Jigsaw Project)

20. Ramaswamy, L., Iyengar, A., Liu, L., Douglis, F.: Automatic detection of fragments in dynamically generated Web pages. In: Proc. of the 13th International World Wide Web Conference (2004). (2004)

21. Naaman, M., Garcia-Molina, H., Paepcke, A.: Evaluation of esi and class-based delta encoding. In: Proc. of the 8th International Workshop on Web Caching and Content Distribution (WCW'03). (2003)

22. Arasu, A., Garcia-Molina, H.: Extracting structured data from Web pages. In: Proc. of ACM SIGMOD'03. (2003)

23. Bar-Yossef, Z., Rajagopalan, S.: Template detection via data mining and its applications. In: Proc. of the 11th International World Wide Web Conference (2002). (2002)

24. Butler, D., Liu., L.: A Fully Automated Object Extraction System for the World Wide Web. In: Proceedings of ICDCS-2001. (2001)

25. Gu, X., et al.: Visual based content understanding towards Web adaptation. In: Proceedings of AH-2002. (2002)