

Evaluation of Edge Caching/Offloading for Dynamic Content Delivery

Chun Yuan, Yu Chen, and Zheng Zhang

Abstract—As dynamic content becomes increasingly dominant, it becomes an important research topic as how the edge resources such as client-side proxies, which are otherwise underutilized for such content, can be put into use. However, it is unclear what will be the best strategy, and the design/deployment trade offs lie therein. In this paper, using one representative e-commerce benchmark, we report our experience of an extensive investigation of different offloading and caching options. Our results point out that, while great benefits can be reached in general, advanced offloading strategies can be overly complex and even counterproductive. In contrast, simple augmentation at proxies to enable fragment caching and page composition achieves most of the benefit without compromising important considerations such as security. We also present Proxy+ architecture which supports such capabilities for existing Web applications with minimal reengineering effort.

Index Terms—Edge caching, offloading, dynamic content, fragment caching, page composition.

1 INTRODUCTION

DYNAMIC pages will dominate the Web of tomorrow. Indeed, one should stop talking about dynamic *pages* but, instead, dynamic *content*. This necessitates architectural change in tandem. In particular, resources that are already deployed near the client such as the proxies that are otherwise underutilized for such content should be employed.

Legitimate strategies include offloading some of the processing to the proxy or simply enhancing its cache abilities to cache fragments of the dynamic pages and perform page composition. While performance benefits including latency and server load reduction are important factors to consider, issues such as engineering complexity as well as security implication are of even higher priority. Although there have been extensive research on the subject of optimizations for dynamic content processing and caching, we still lack the insight on what will be the best offloading and caching strategies and their design/deployment trade offs.

In this paper, using a representative e-commerce benchmark, we have extensively studied many partitioning strategies. We found that offloading and caching at edge proxy servers achieves significant advantages without pulling databases out near the client. Our results show that, under typical user browsing patterns and network conditions, two to three folds of latency reduction can be achieved. Furthermore, more than 70 percent of server requests are filtered at the proxies, resulting in significant server load reduction. Interestingly, this benefit can be achieved largely by simply caching dynamic page fragments and composing the page at the proxy. In fact, advanced offloading strategies can be overly complex and

even counterproductive performance wise if not done carefully. Our investigation essentially boils down to one simple recommendation: If end-to-end security is in place for a particular application, then offload all the way up to the database; otherwise, augment the proxy with page fragmentation caching and page composition. While our results are obtained under the .NET framework, we believe they are generic enough to be applicable to other platforms.

The rest of the paper is organized as follows: Section 2 covers related work. Various offloading and caching options are introduced in Section 3, which also discusses important design metrics. Section 4 examines the benchmark used and also some of the most important .NET features employed. Detailed implementations of the offloading/caching options are discussed in Section 5. Section 6 describes the experiment environment. Results and analysis are offered in Section 7. Section 8 describes Proxy+ architecture which supports dynamic content caching with minimal engineering overhead for existing applications. Finally, we summarize and conclude in Section 9.

2 RELATED WORK

Optimizing dynamic content generation and delivery has been widely studied. The main objectives are to reduce client response time, network traffic, and server load caused by surges of high volume of requests over wide-area links. Most work focuses on how to support dynamic content caching on server side [9], [10], [19]. Some others also extend their cache to the network edge and show better performance results [11]. Fragment caching [3], [4] is an effective technique to accelerate current Web applications which usually generate heterogeneous content with complex layout. It is provided by today's common application server product like Microsoft ASP.NET [14] and IBM WebSphere Application Server [7]. ESI [5] proposes to cache fragments at the CDN stations to further reduce network traffic and response time.

• The authors are with Microsoft Research Asia, 5F Sigma Center, #49 Zhichun Road, Beijing, 100080, China 86-10-62617711.
E-mail: {cyuan, ychen, zzhang}@microsoft.com.

Manuscript received 26 Oct. 2003; revised 1 Apr. 2004; accepted 8 Apr. 2004.
For information on obtaining reprints of this article, please send e-mail to: tkde@computer.org, and reference IEEECS Log Number TKDE-0203-1003.

Application offloading is another way to improve performance. In Active Cache [2], it is proposed that a piece of code be associated with a resource and be cacheable, too. The cache will execute the code over the cached object on behalf of the server and return the result to the client directly when the object is requested at a later time. With the blurring of application and data on current Web, this scheme becomes less effective since Web page generation and delivery involve the most overhead and a Web page is typically constructed from specific data sources rather than from a Web resource. To do more aggressive application offloading, WebSphere Edge Services Architecture [8] suggests that portions of the application such as presentation tier and business logic tier be pushed to the edge server and communicate with the remaining application at the origin server when necessary via the application offload runtime engine. Gao et al. [6] proposed to take advantage of application semantics to replicate some data objects to the edge server and observed significant performance improvement while maintaining reasonable consistency. This is an instance of pushing database tier (partially) to the edge, though in an opaque way. An extreme case of offloading is given by [1]. The full application is replicated on the edge server and database accesses are handled by a data cache which can cache query results and fulfill subsequent queries by means of query containment analysis without going to the backend.

We focus on the proxies that are already installed near clients. We also examine exclusively offloading and caching of anything other than the database content, as we believe mature technologies to manage hard states in a scalable fashion across wide-area are yet to be developed. To the best of our knowledge, we are the first to report design and implementation trade offs involved in devising partitioning and offloading strategies, along with detailed evaluations. There also has been no work evaluating offloading versus advanced caching mechanisms. Finally, this is the first work we know of that experiments with the .NET framework in this aspect.

3 OFFLOADING AND CACHING OPTIONS ENUMERATED

There are a number of issues to be considered for distributing, offloading, and caching dynamic content processing and delivery. They are:

1. available resources and their characteristics,
2. the nature of these applications, and
3. a set of design criteria and guidelines.

In this section, we discuss these issues in turn.

3.1 Resources Where Offload can be Done

Fig. 1 shows graphically various resources involved.

Client. As a user-side agent, client—typically a browser—is responsible for some of the presentation tasks and it can also cache some static contents such as images, logos, etc. The number of clients is potentially many; however, they usually have limited capacities and are (generally speaking) not trusted.

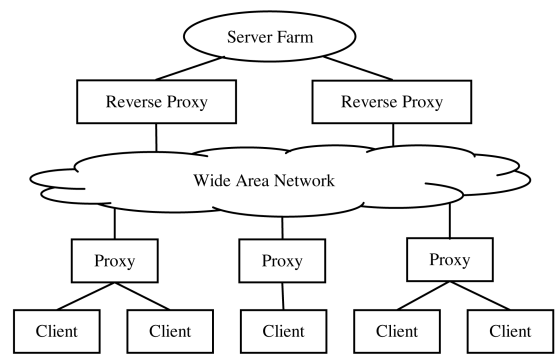


Fig. 1. Resources available for offloading and caching.

Proxies. In terms of scale, proxies come second. Proxies are placed near the clients and are thus far from the server end. The typical functionalities of proxies include firewall, and caching of static contents. They are usually shared by many clients and are reasonably powerful and stable. However, except the case of intranet applications, content providers do not have much control over them.

Reverse Proxies. Reverse proxies are placed near the back-end server farm and act as an agent of the application provider. They serve the Web request on behalf of the back-end servers. Content providers can fully control their behaviors. However, the scale of reverse proxies only goes as far as a content provider's network bandwidth allows. In this paper, we consider them as part of the server farm.

Server. Servers are where the content provider has the full control. In the context of this section, we speak of "server" as one logical entity. However, as it shall be clear later, "server" itself is a tiered architecture comprised of many machines and hosting the various tiers of the Web application.

As far as dynamic content is concerned, typically, only the servers and clients are involved. Proxies, as of today, are incapable of caching and processing dynamic content. In this discussion, we have also omitted CDN stations as we believe they can be logically considered as an extension of either proxies or reverse proxies. Some of the more recent progresses have been discussed in the Section 2.

3.2 Application Architecture and Offloading Options

Logically, most of the Web applications can be roughly partitioned into three tiers: presentation, business logic, and back-end database. The presentation tier collects users' input and generates Web pages to display results. The business logic tier is in charge of performing the business procedure to complete users' requests. The database tier usually manages the application data in a relational database.

Based on the three-tier architecture, N-tier architecture is also possible. The most complex tier in a Web application is the business logic tier. This tier performs application-specific processing and enforces business rules and policies. Because of its complexity, the business logic tier itself may be partitioned into smaller tiers, evolving into the N-tier architecture.

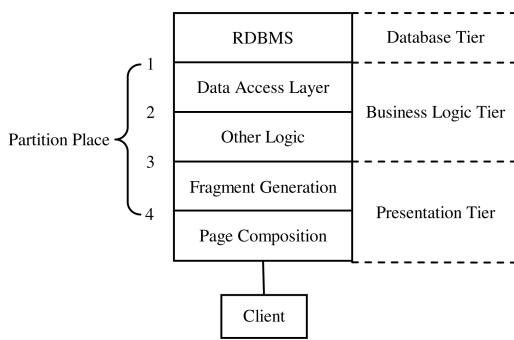


Fig. 2. The 3-tier architecture and partition places.

Application partitioning and offloading can be applied based on the tier structure of the Web application. Without loss of generality, we only consider the Web browser as the application client here. From the back-end database to browsers, we can find several candidate partition places as shown in Fig. 2.

The first partition place is the database access interface. The ODBC, JDBC, ADO, etc., are this kind of interface. Current applications use connection strings to specify the database server to be accessed. It is possible to point to a specific remote machine in the connection string.

The second partition place is at the data access layer inside the business logic tier. Because of the complexity of the business logic, it is a common practice to develop a set of database access objects to shield the detail inside the database. Other business logic objects can access data through these objects using simple function calls. The clear-cut boundary at this layer makes it a good candidate of partition point for offloading.

The third partition place is between the presentation tier and the business logic tier. The presentation tier gathers the user input and translates the user request into a processing action at the business layer. The business tier usually provides a single-call interface for each type of requests. The clearly defined interface here provides strong clues for partition.

The fourth partition place is inside the presentation tier. The Web pages generated by the application are structuralized and split into fragments, each of which has consistent semantic meaning and lifetime. The back-end servers provide page fragments and composition frameworks. The entire Web page is assembled at the offloading destination. ESI [5] is a good example of this strategy.

Of course, what we have enumerated here is only a starting point. Specifically, within the business logic tier there can be multiple logically legitimate offloading points. However, as we should discuss later, advanced offloading strategies often risk high complexity without a clear benefit in return.

3.3 Important Factors to be Considered When Offloading

Having discussed various resources upon which offloading and caching can be performed, and various partitioning strategies, the actual implementation and deployment must consider a number of important factors. In our opinion, the

following three are the most important ones: *security*, *complexity*, and *performance*:

1. **Security.** Sensitivity of data as well as processing that are to be offloaded may vary. A given piece of data and processing can be distributed as far as its security perimeter permits. This is one reason why we are concerned with who controls what in the resource distribution earlier. Enforcing security end-to-end only applies to certain Web applications (e.g., intranet) and pays a cost (e.g., VPN overhead) in return.
2. **Complexity.** Another factor that should be considered is the engineering cost. Although Web applications are developed according to three-tier or N-tier architecture, the tier boundaries are usually not clear. This problem is obvious for the tiers that are part of the business logic in an N-tier application. Even if the tier boundaries are clear, the implementation still cannot be fully automated. For example, application partitioning usually requires transforming some of the LPCs (local procedure call) into RPCs (remote procedure call). Because most of the runtime systems do not support migrating LPC to RPC transparently, source code modification, recompilation, and subsequent testing are necessary. If synchronous procedures calls are to be changed to asynchronous calls, the implementation efforts would be even greater.
3. **Performance.** Even when resources such as proxies are freely available, distributing the processing and caching must bring significant benefits to justify the additional complexity involved. End user's latency as well as improvement of scalability are the primary metrics. On this, the network condition is the first critical factor to be considered. Generally speaking, the communication quantity across the partition should be minimized on low bandwidth networks. Likewise, for high latency networks, the frequency of synchronous communication should be reduced. In general, a useful guideline to start with is that the communication channel over a wide area network should be lightweight and stateless.

4 THE PET SHOP BENCHMARK

In order to evaluate different offloading options, we use Microsoft .NET Pet Shop as our benchmark. It comes from Sun's primary J2EE blueprint application, the Sun Java Pet Store [18] and models a typical e-commerce application, an online pet store. E-commerce sites like this are among the most common Web applications.

Pet Shop is implemented using ASP.NET, and the source code is freely available at [12]. ASP.NET brings several important optimizations and the two of them, *stored procedure* and *output caching*, will be discussed in the following sections.

4.1 Pet Shop Architecture

The complete three-tier architecture of Pet Shop is described in the white paper at [12]. To illustrate the design, we will

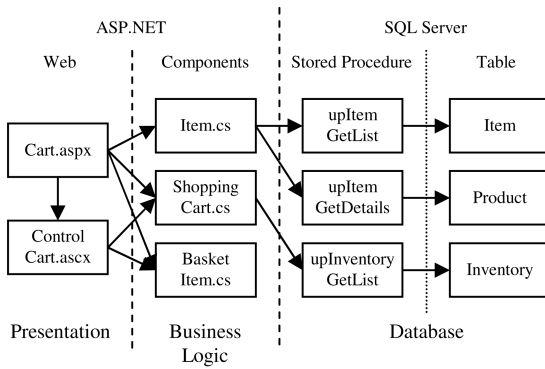


Fig. 3. The Pet Shop architecture (portion).

look at an example of the interaction between the three tiers as shown in Fig. 3.

The *presentation tier* communicates with browsers directly. It contains *Web Forms pages* (*aspx* files), *Web Forms user controls* (*ascx* files), and their *code-behind* classes (in namespace *PetShop.Web*). Similar to the ASP and JSP page, Web Forms pages represent dynamic pages. The Web Forms user controls represent portions of Web Forms pages and, thus, cannot be requested independently. While the *aspx* and *ascx* files contain the visual representation, the code-behind classes contain processing logic. When a request arrives, the specified Web Forms page and Web Forms user controls are loaded. The corresponding code-behind objects responsible for generating responses will initiate calls to the business logic tier for request processing (arrows in Fig. 3).

The objects in the *business logic tier* (in namespace *PetShop.Components*) accept invocations from the presentation tier. If the processing does not require database interaction, for instance, displaying shopping cart content, results are returned right away. Otherwise, the business logic objects will generate database queries through a specific database access class (*PetShop.Components.Database*). Instances of this class set up database connections, pass database queries through ADO.NET interfaces, and return query results to upstream.

The *database tier* consists of application data and *stored procedures*. A *stored procedure* is used to encapsulate a sequence of SQL queries which complete a single task. Using *stored procedures*, interactions between the business logic tier and the database tier can be reduced, thus increasing performance. For instance, placing an order normally requires several calls between the business tier and the back-end database. With stored procedure, an order can be encoded into a string and transferred to the database, where the string is decoded and multiple SQL statements are issued to complete the order. From this perspective, most of the Pet Shop *stored procedures* are essentially part of the business logic tier. They are included in the database tier simply because they are stored and are executed in the SQL server. This is one example where the boundaries of tiers get blurred.

4.2 ASP.NET Output Caching

The.NET Pet Shop leverages ASP.NET output caching to increase throughput and reduce server load [13]. A similar




Fig. 4. Implementation of offloading options in Pet Shop.

function is also provided by other products such as IBM WebSphere's *response cache* [7], [8]. When a page is requested repeatedly, the output caching allows subsequent requests to be satisfied from the cache so the code that initially creates the page does not have to be run again. Besides caching the entire page, ASP.NET allows Web Forms user controls to be cached separately. As we will explain in detail in Section 5.3, this feature of fragment caching is what we employ to enhance caching capability at the proxy side.

ASP.NET provides duration and versioning control for each cached entity (*Web Form* and *Web Form user control*). Duration specifies the lifetime of a cached page. Versioning allows caching multiple result pages or page fragments for a single form or control. For example, *Product.aspx* produces different result pages for different products. Storing a single result page in output cache can hardly gain any benefit since users tend to browse different products. By keeping multiple result pages, the most frequently accessed pages will be cached eventually, saving large amounts of processing time.

5 EXPERIMENT PREPARATION

In this section, we discuss in detail how different offloading and caching strategies are implemented in Pet Shop.

According to the partition points in Section 3.3, the following offloading options are investigated: F_0 , F_{db} , $F_{remoting1}$, and F_{proxy} . They are shown in Fig. 4; the legends are:

- B: Browser,
- A: Page Assembling and Fragment Caching,
- G: Fragment Generation,
- P: Presentation,
- L': Business Logic except Data Access Layer,
- DA: Data Access Layer,
- L: Business Logic,
- DB: Database, and
- Cloud: Wide Area Network.

The base line is F_0 which leverages neither processing nor caching abilities of proxies. By pushing fragment caching and page assembly to the proxy, we get F_{proxy} which corresponds to partition place 4. By offloading the presentation tier to proxies, $F_{remoting2}$ implements partition point 3. $F_{remoting1}$ and F_{db} are similar except that $F_{remoting1}$ leaves the data access layer at back-end servers while F_{db}

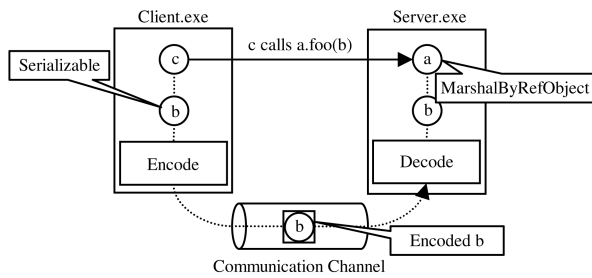


Fig. 5. RPC using .NET Remoting mechanism.

offloads the complete business logic tier. Therefore, they correspond partitioning points 2 and 1, respectively.

We use proxy and front end, and server and back end interchangeably in this paper for all configurations other than F_0 .

5.1 Implementing F_{db}

The implementation of F_{db} is trivial: All we need to do is modify the connection string of the database access interface. The connection string is changed from the default value “server = localhost; ...” to “server = some other machine; ...” so that the front end is forced to access the remote machine hosting the SQL server.

5.2 Implementing $F_{remoting}$

The $F_{remoting}$ option investigates different ways of offloading inside the business logic tier, in particular, the partition points 2 and 3 (see Fig. 2). We employ the .NET Remoting feature to accomplish this task, which we will discuss first.

5.2.1 .NET Remoting

Microsoft .NET Remoting provides a rich and extensible framework for objects living in different application domains, in different processes and in different machines to communicate with each other seamlessly. The framework considers a number of matters, including *object passing*, *remote object hosting strategy*, *communication channel*, and *data encoding*.

Objects can be passed either by reference or by value:

- **By value.** Objects that would cross application domain boundary, such as the object *b* in Fig. 5, can be passed by value. In .NET Remoting, pass-by-value objects are all marked with *Serializable* attribute.
- **By reference.** Objects that reside in only one application domain and provide interfaces to other applications are passed by reference (such as the object *a* in Fig. 5). In .NET Remoting, pass-by-reference objects should be derived from a system class *MarshalByRefObject*.

For a pass-by-reference object, .NET Remoting provides three hosting strategies to support object activation and lifetime management:

- *SingleCall* objects’ activation and lifetime are determined by server. They service one and only one

request coming in, i.e., different client requests are services by different objects.

- *Singleton* objects’ activation and lifetime are also determined by server. Unlike the *SingleCall* objects, *Singleton* objects service multiple clients and share data by storing state information between client invocations. There is only one singleton object instance of a given class at the server side.

Client-activated objects’ (CAO) activation and lifetime are determined by client. The server creates an object upon an activation message from client. The object services for the client until the client allows it to be released. If the communication between server and client is stateful, CAO should be used.

For a more in-depth treatment of these hosting strategies, please refer to [16].

RPC requests and responses are encoded into formatted messages and transferred over a communication channel. In Fig. 5, when object *c* makes a call to object *a*, the request (including object *b* as the parameter) is encoded and transferred to the server side. At the server side, the message is decoded and an actual call to object *a* is made.

5.2.2 Detailed Implementation

As explained earlier, in $F_{remoting}$, we try to partition the application in the logic tier. While there may be many different options, as an extensive exercise, we investigate how to partition at point 2 which separates the data access layer from other business logic layers, and point 3 which is located between the presentation tier and the business logic tier (see Figs. 2 and 4).

Regardless of the specific partition strategy, the task in this configuration is always to replace LPC (Local Procedure Call) with RPC (Remote Procedure Call). This entails a few steps. The first is to determine the locations of classes to be run as mandated by a given partitioning strategy (server or proxy) and from there derive the RPC boundaries. The second step is to modify the application source code so that RPC can take effect. Finally, the hosting strategy for objects at server side and communication channel between server and proxy are decided.

Partition place 2 requires the least amount of engineering efforts in that there is only one class to be modified—Database in namespace *PetShop.Components*. The Database objects run at server side and provide interfaces for the other logical tier objects to access information in the back-end database. Thus, they are pass-by-reference objects. For each user request, the responsible logic tier object at proxy side needs to issue multiple procedure calls to server (as shown in Fig. 6a). Because these calls are related to each other, the state information along the call sequence should be maintained. On the other hand, different user requests need exclusive objects to provide services. Therefore, the only hosting strategy is CAO. This strategy turns out to have a dramatic performance impact: Our test runs reveal that the benchmark now performs much worse than not offloading *at all*. The reason is that multiple RPCs corresponding to a single-user request results in multiple round trips between proxy and server. Consequently, partition point 2 is not a good offloading option for Pet

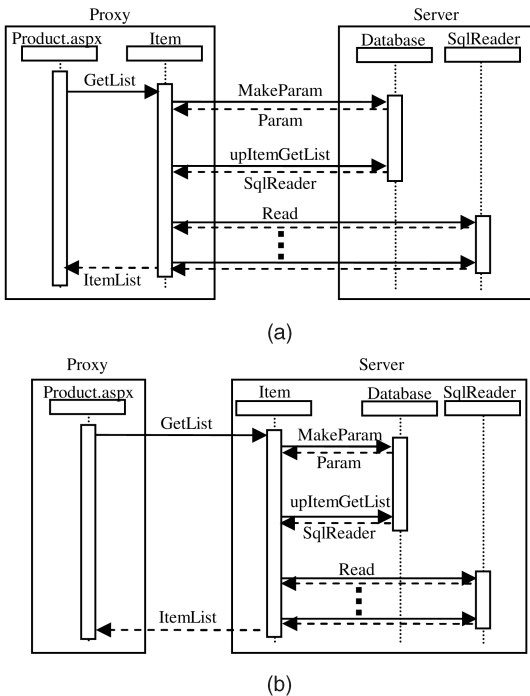


Fig. 6. Flow examples of partition place (a) 2 and (b) 3.

Shop. The lesson we learned here is that, under a given partition strategy, the logical constraint and performance constraint may be very much in conflict with each other. What makes this matter even more complex is that this is also application dependent.

Partition point 3 works but with a rather involved manual examination of all candidates. Specifically, all classes in namespace PetShop. Components except Error and Database are affected.¹ Some of them will reside at server side only and act as pass-by-reference remote objects, such as Item, Profile, Order, Customer, and Profile. The others are pass-by-value objects that will travel between proxy and server, such as ShoppingCart,² BasketItem, ItemResults, ProductResults, and SearchResults (actually, they are all RPC parameters or responses). With this strategy, the proxy needs to issue only one RPC for each user request in most cases and no state information is going to be shared among different RPCs (as shown in Fig. 6b). Moreover, because the objects of Item, Profile, Order, Customer, and Profile do not store state data,³ there is no difference whether distinct user requests are processed by a shared object or exclusive and different ones. Therefore, all the three hosting strategies are eligible. We conducted an experiment to compare the performance of all of them under the same test condition (as shown in Table 1) and found that *SingleCall* and *Singleton* outperforms *CAO* significantly. In the experiment, we notice that, under the same conditions, the memory consumption and CPU load on server for *CAO* is higher than that for *SingleCall* and *Singleton*. This may be due to

1. Error objects are responsible for reporting local errors. Although Database objects reside at the server side, they do not provide interfaces for proxy any more in this option.

2. Actually, ShoppingCart will access database in Cart.aspx when updating the shopping cart information. To deal with it, we modify Cart.aspx and Order to ensure that the update would be executed on server.

3. In the C# source code, these five classes have no member variables.

TABLE 1
Comparison of *Singleton*, *SingleCall*, and *CAO*

| | Requests/sec | Response time (ms) |
|------------|--------------|--------------------|
| Singleton | 649.07 | 24.66 |
| SingleCall | 649.17 | 24.46 |
| CAO | 552.31 | 43.75 |




Fig. 7. Page fragment composition.

the fact that the server no longer can perform aggressive garbage collection as it can in the other two options.

In the following sections, $F_{remoting}$ will refer to $F_{remoting2}$ using *Singleton* only. While we did arrive at an adequate partition strategy in the business logic tier for Pet Shop (see performance results in Section 7), our experience pointed out that this is a rather complex process and, even though it is possible to derive a consistent set of guidelines, it will be quite a challenge to perform automatic partitioning.

5.3 Implementing F_{proxy}

As we described earlier, the goal of F_{proxy} is much more modest: Augment the proxy cache so that it can cache fragments and perform dynamic page assembly. Recall that the objects in the *presentation tier* of the original Pet Shop are divided into two parts: container pages (*Web Forms*) and fragments (*Web Form user controls*). Each container page includes placeholders of some fragments.⁴ Their contents, either obtained from the output cache or generated afresh upon a miss, are to be inserted into the container at runtime to compose a complete Web page.

The flow of F_{proxy} is shown in Fig. 7. In order to accomplish this task, we need to 1) replicate the tier responsible for output caching and page assembly at the front end and 2) make a back-end version of Pet Shop which is the real generator of content.

The back-end version is implemented in two steps. First, we make each fragment be able to be separately retrieved with a URL like a common Web page. Then, in their container pages, we replace the fragment's placeholder with a special tag that will not be interpreted by ASP.NET, but indicate to the front end that it is to be expanded with the content of a fragment.

We create another application to play the role of output caching and page assembly at the front end. This application has the same set of pages and fragments as Pet Shop and their containment relationships are also maintained, but no actual content is included. When a page is loaded, the container page and the nested fragments will be loaded in turn. What they actually do

4. The container object of a fragment can be a fragment, too. In Pet Shop, all fragments are contained in a page.

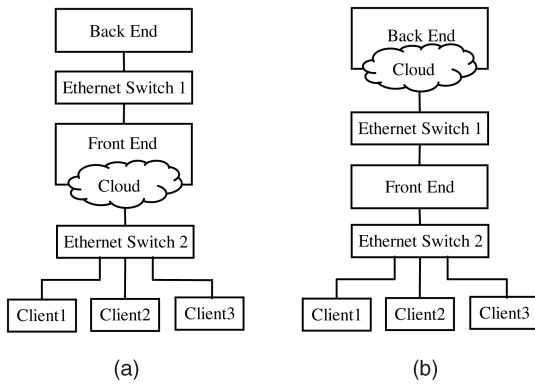


Fig. 8. Network configuration in the experiment.

is retrieving their corresponding page” from the output cache in case of cache hits, or from the back-end application, otherwise. After the real content has arrived, the page composition begins. The special tags in the fetched container page are replaced with the content of its subordinate fragments. The process will be recursively performed if there are nested containments, until the page is finally composed.

In summary, here are the steps that will occur at runtime:

1. client sends request to proxy,
2. proxy fetches container and fragments either from the output cache (now hosted at proxy) or request from backend, and
3. proxy assemble the page and returns to the client.

We should point out that this version of implementation is for a quick evaluation of the F_{proxy} option and is much more like a hack: There are no modifications to either proxy or .NET framework anywhere, we simply replicate some of the .NET functionality at the proxy side. A more solid and application-independent implementation will be described in Section 8.

6 EXPERIMENT SETUP

6.1 The Test Bed

There are a total of five machines in our experiment: two servers and three clients (as shown in Fig. 8). Both servers are powerful machines each with two Pentium 4 2.2GHz CPUs, 2GB RAM, and two 73GB SCSI disks. The clients are PCs with sufficient power so that they never become bottlenecks in test runs. The three clients are connected to the front-end server through a 100Mbps Ethernet switch. In order to avoid contentions between front-end—client communication and front-end—back-end communication, two network adapters are installed on the front-end server. The two portions of traffic will go through different network adapters and switches.

In the test runs, F_0 uses Fig. 8a, while the other three configurations use Fig. 8b:

- For F_0 , the front-end server runs the original Pet Shop application and the back-end server runs database.
- For $F_{remoting}$, the presentation tier is hosted at the front end and all the rest, including the business logic tier and the database, are at the back end.

TABLE 2
Software Configuration in the Experiment

| | |
|------------------|---|
| Web Server | IIS 5.0 ¹ + ASP.NET ² |
| DBMS | SQL Server 2000 |
| Operating System | Windows 2000 Advance Server SP 3 |
| Network Emulator | Shunra\Cloud |

1. In order to obtain reasonable performance, the application protection mode is set to medium.

2. Come with .NET Framework V.1.0.3705.

- For F_{db} , the original Pet Shop runs at the front end which accesses the database at the back end server.
- For F_{proxy} , the page assembly and fragment caching tier runs at the front end and the modified Pet Shop runs at the back end accessing a local database.

The software environment of our test bed is shown in Table 2.

The Shunra\Cloud (version 3.1) [17] is used to emulate network latency in our experiment. It is attached to a network adapter and affects all the IP packets through it. In Fig. 8a, Cloud resides in the front-end server to emulate WAN conditions between clients and Web server. In Fig. 8b, Cloud is associated with the back-end server to emulate WAN conditions between proxies and content provider’s servers. Cloud only imposes minor additional loads on the server it attaches and its effect is negligible.

6.2 Client Emulation

We use Microsoft Application Center Test (ACT) to emulate surges of clients. For each test, ACT distributes test load to the client machines. Each client creates enough threads to simulate a number of concurrent Web browsers visiting the Web application under test. The actual behavior of the threads is controlled by a test script. In each test, a thread repeats the following steps until the test duration is over. It first opens a persistent HTTP connection to the Web server. Then, it chooses a request and sends it to the Web server. After receiving the response, it waits for some thinking time (50 milliseconds in our tests) and chooses the next request and so forth. The selection of the request to send is determined by a state transition diagram, which defines the probability of going to the next request from the previous one. If the thread chooses to exit, it stops sending requests and closes the connection.

TABLE 3
Distribution of the Test Workload

| Activity | Percentage |
|------------------------|------------|
| Category Browsing | 18% |
| Product Detail | 16% |
| Search | 18% |
| Home Page | 18% |
| Shopping Cart | 7% |
| Order | 1% |
| Account/Authentication | 22% |

The workload generated by the test script roughly follows the distribution shown in Table 3, which corresponds to typical user browsing patterns for such Web site.

6.3 Tuning

Through our experiments, we found it is necessary to finetune the configurations so as to eliminate as many side-effects as possible. The major tunings are reported in this section.

6.3.1 Queuing and Threading

ASP.NET assigns processing and I/O jobs (via .NET Common Language Runtime (CLR)) to worker threads and I/O threads in a thread pool with a specified maximum size to process the requests from an incoming queue for the application requested. For some configurations in our experiment, satisfying a request at the front-end server may require accessing external resources (i.e., object, database, page fragment) from the back-end server with a network delay of hundreds of milliseconds. The worker thread has to be blocked and wait for I/O to complete. Therefore, a high number of concurrent requests will lead to many blocked threads in the pool as well as pending requests in the queue. To minimize this kind of effect, we make the following adjustments:

1. We set every application's request queue long enough to guarantee that requests be normally served instead of being rejected with an HTTP message indicating server error when server-side congestion occurs. In reality, Web servers limit the length of request queues (for dynamic page) to prevent lengthy response time that is not acceptable to users and, therefore, save resources for incoming requests. In our experiment, we need to measure the normal response time for each request, so we raise the default queue length limit from 100 to 1,000.
2. Thread pool size affects performance significantly in a subtle way. Too few threads can render the system underutilized because there may be many blocked threads. While increasing thread pool size might accelerate processing and improve utilization, it will cause more thread context switch overhead. In our experiment, we tune the thread pool size at the front-end server for each configuration under each test condition (various network latencies and output cache hit ratio) to produce the optimal result (throughput and utilization). In general, a configuration with a longer processing time caused by network latency needs a larger thread pool size. For example, when the network latency is 200ms and the output cache hit ratio is 71 percent, the optimal thread pool size of F_0 , F_{db} , $F_{remoting}$, and F_{proxy} is 25, 30, 30, and 50, respectively.

6.3.2 Connection Pooling

In the three configurations other than F_0 , the front-end server needs to make heavy communications with the back-end server via TCP. On a network with high latency, the connection overhead becomes prominent due to TCP handshake and slow start effect if no connection pooling

is used. In order to prevent this kind of performance degradation, we set appropriate pooling parameters for each configuration according to the connection mechanisms used between the front end and the back end.

In F_{db} configuration, each time the front end wants to create a connection to the back-end database, the underlying data access component (ADO.NET) will pick a usable matching connection from a pool of a certain size. If there is no usable connection and the size limit is not reached, a new connection will be created. Otherwise, the request is queued before a timeout error occurs. When the network latency is high and the work load is heavy, the number of concurrent connections in use would constantly exceed the limit, i.e., the connection pool becomes the bottleneck. Therefore, we set the connection pool size large enough (300 rather than the default 100).

In $F_{remoting}$ configuration, front-end objects invoke back-end objects through .NET Remoting TcpChannels, which will open as many as needed connections and cache them for later use before closing them after 15-20 seconds of inactivity. Since each test run has a warm-up period, there is no negative impact on performance from this setting.

In F_{proxy} configuration, the front end needs to retrieve Web objects from the back end. We set the maximum number of persistent connections high enough (1,000) to avoid congestions.

6.4 Measurement

Our experiment consists of two parts. The first part measures the performance of the four configurations under three fixed network latencies (50ms, 200ms, and 400ms). The second part measures the performance under several cache hit ratios.

For each test configuration, we vary the number of concurrent connections and run ACT to stress the application and measure the throughput (in terms of requests per second or RPS), response time (in terms of the time to the last byte of a response), and utilization of the servers (front end and back end). Each run starts with a warm-up period (3 minutes), after which ACT begins to collect data about the system status every one second by reading performance counters of processors, memory, network interfaces, ASP.NET, .NET CLR, and SQL Server. Data collection traffic only consumes a slight portion of bandwidth compared to client requests and server responses so it interferes little with the network utilization of test load. The collection process will last for 3 minutes and then ACT stops running. Before each test run, we reset the Web server (to flush the cache) and the database tables (to restore the data load) so that the results from different runs are independent.

7 RESULTS AND EVALUATION

In this section, we offer detailed experiment results and analysis. Due to the large number of configuration combinations, we cannot present all the data. Thus, we report response time and resource utilization with one representative network latency in Section 7.1, and briefly go over more advanced variations in Section 7.2.




Fig. 9. Response time versus number of connections.

7.1 Basic Results

7.1.1 Response Time

We find F_{proxy} yields response time comparable to that of $F_{remoting}$ and F_{db} , which is partly due to its efficiency of communication between the front end and the back end. To further clarify the result, we also show performance figures for three kinds of requests that constitute the work load below.

Fig. 9 shows the average response time versus concurrent connections number of the four configurations. The cache hit ratio is 71 percent, and the network latency is 200ms. The loads on the front-end server in the three offloading options are higher relative to that of F_0 , either due to more processing or higher overhead. This is the reason for the climbing of the latency curves corresponding to the three optimizations. This is an artifact of our test bed where there is only one proxy and is generally not a concern because proxies are many in real deployment.

When there are less than 300 connections, F_{db} , $F_{remoting}$, and F_{proxy} offer a response time better than F_0 . The reduction depends on the network delay. When the number of connections is 15, the reduction is 76, 75, and 64 percent, for F_{db} , $F_{remoting}$, and F_{proxy} , respectively. In F_0 , every time a request is issued from the client, it travels through the delayed link and so does its response. Therefore, the response time is always above the network roundtrip time (400ms). In the other three configurations, many requests can be satisfied at the front end right away. F_{db} and $F_{remoting}$ achieve this by replicating the application logic fully or partially from the back end. F_{proxy} caches dynamically generated output such as pages and fragments and serves subsequent matching requests from the cache directly or by composing fragments in the cache, thus reducing the frequency of accessing the back end significantly.

The response time of F_{db} and $F_{remoting}$ are very close under small number of connections. This is because their partition points afford them the same back-end access rate and the traffic incurred by each access is enough to be transferred within one roundtrip for both configurations.

Compared to F_{proxy} , F_{db} and $F_{remoting}$ do more than just caching: They are capable of offloading some of logical processing as well. For example, the responses of some




Fig. 10. Average response time of each class of pages.

requests, such as `GET CreateNewAccount.aspx` and `GET SignIn.aspx`, are not cacheable, so they will always go to the back end in F_{proxy} . But, neither F_{db} nor $F_{remoting}$ requires back end access for these requests.

The fact that the cacheable pages occupy a large portion of the test loads is one reason that F_{proxy} does not lag too far behind. However, there is one more subtle and interesting case where F_{proxy} actually wins out.

F_{proxy} is optimized for retrieving Web pages and fragments from the back end and performing page assembly. When it encounters a request for a page containing fragments, all objects including the container page and the fragments inside that are not cached will be fetched from the back end asynchronously, allowing them to be downloaded in parallel. After all objects arrive (in any order), they are composed together and a complete page will be returned. While in F_{db} or $F_{remoting}$, due to the fact that the partition points are inside the application logic, all accesses to the back end such as database query and remote object invocation are blocking. For example in $F_{remoting}$, if completing a request needs to call the back end twice for generating two fragments in the page, the two invocations must be performed sequentially even though the two fragments are independent of each other. The asynchronous retrieving nature of F_{proxy} saves significant time for processing such pages compared to F_{db} and $F_{remoting}$. For instance, the page `Cart.aspx` (uncacheable, used for adding, updating, and removing items in a shopping cart) contains one cacheable fragment and two uncacheable ones which show a favorite list and a banner, respectively. The favorite list and the banner need to query the back-end database and the container page also needs to in case of adding an item to the shopping cart. When the network latency is 200ms and the work load is light it takes F_{db} and $F_{remoting}$ 1,245ms and 1,285ms, respectively, (both over three roundtrips), but takes F_{proxy} 784ms to generate the page.

Therefore, we can classify requests into three distinctive classes (with their ratios):

- Class A (71 percent): cacheable response,
- Class B (14 percent): uncacheable response and may need to access database, and
- Class C (15 percent): also uncacheable response but there is no need to access database.

Fig. 10 compares the average response time of the requests in each class for $F_{remoting}$, F_{db} , and F_{proxy} . All of

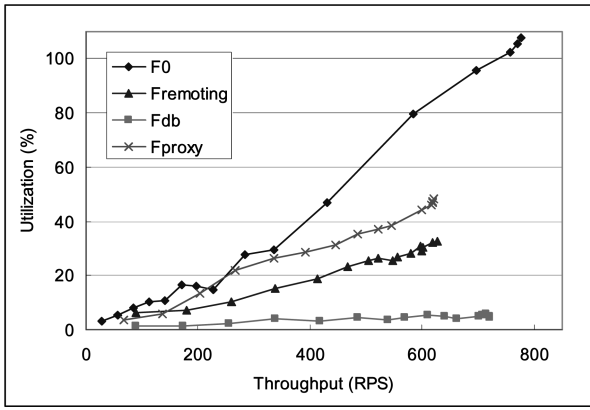


Fig. 11. CPU utilization in several of all configurations. **Note: The three curves of the offload configurations stop at throughput points beyond which our front end proxy becomes saturated.**

them return responses for requests of class A in negligible time. F_{proxy} wins in class B due to its asynchronous optimization mentioned above. $F_{remoting}$ responds slower than F_{db} because it introduces overhead when doing remote object invocations. In class C, both F_{db} and $F_{remoting}$ can return responses immediately, while F_{proxy} has to fetch content from the back end with significant delay.

7.1.2 Scalability and Server Load Reduction

The other functionality of offloading and caching at proxies is reducing server load so as to achieve better scalability. Since our test bed consists of only one proxy and one database server, we can only infer from the load distribution among resources. However, all offloading configurations filter at least 70 percent of server requests through proxy due to the output cache hits.

Fig. 11 plots the aggregated server loads for all four configurations. These are loads that will still remain at the server side regardless in real deployment. The curve of F_0 adds up loads of both the front end (which runs the application) and the back end (the database). Curves of all three offload configurations are the loads on the back-end server only and they are:

- F_{db} : the database,
- $F_{remoting}$: the remaining of application logic and the database, and
- F_{proxy} : modified Pet Shop application and the database.

As we explained earlier, the front-end server in these configurations all have higher loads than in F_0 , either because of more processing or higher overhead. Consequently, the three curves of the offload configurations do not extend to as high throughput as in F_0 because our front end proxy becomes bottleneck.

As can be seen, in our test environment, the back-end database is not the bottleneck. The largest load reduction is achieved with F_{db} where the functions of application server are taken entirely by the proxies distributed near the client. Because there is still some processing remaining at the back end in $F_{remoting}$, it cannot achieve the same level of server

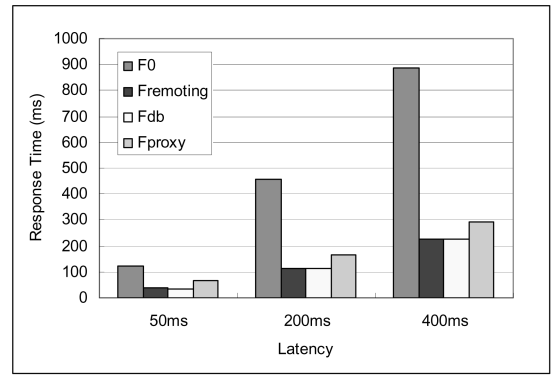


Fig. 12. Response time with different network latencies.

load reduction as F_{db} . As expected, F_{proxy} is the third and achieves reasonable load reduction comparing to F_0 .

In reality, a server complex is made up of a tier of machines running application servers backed by the database machines. The differences between the load curves of F_0 , $F_{remoting}$, F_{proxy} , and F_{db} are what will be run on the machines hosting the application servers. It is evident that to achieve identical throughput, F_{db} will require the least resources inside server complex, followed by $F_{remoting}$ and F_{proxy} .

7.2 Advanced Evaluation

7.2.1 Vary Network Latency

We repeat the experiment under two other network latencies: 50ms and 400ms with light server load (15 connections). The results are shown in Fig. 12.

There is no major surprise: The response time of F_0 is always over one roundtrip time. Both F_{db} and $F_{remoting}$ offer comparable response time and F_{proxy} is slower but still competitive. As expected, the benefit of offloading/caching at proxy increases with network latency.

7.2.2 Vary Cache Hit Ratio

We repeat the experiment under two other output cache hit ratios: 52 and 30 percent. The network latency is fixed at 200ms. The response time of the configurations under different hit ratios is shown in Fig. 13.

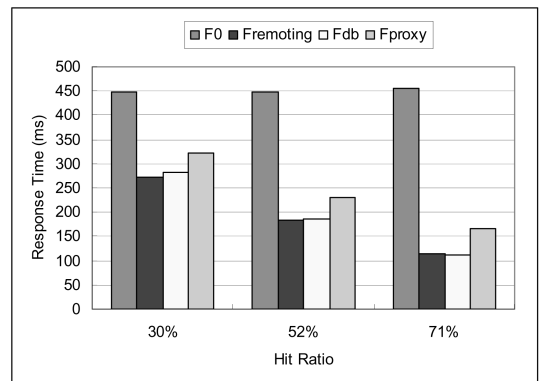


Fig. 13. Response time of different cache hit ratios.

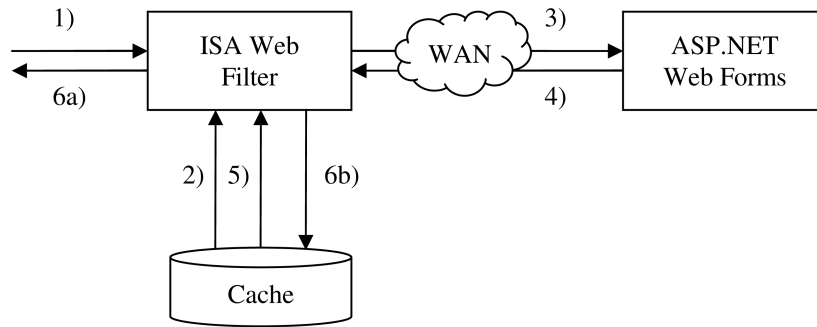


Fig. 14. Architecture and workflow of Proxy+.

This result is also straightforward: Increased cache hit ratio benefits offloading/caching at proxy. We note again that the caching is a more significant factor, though this is specific to this application and the test scripts we used for the experiments.

8 PROXY+: SIMPLE PROXY AUGMENTATION FOR DYNAMIC CONTENT PROCESSING

As stated before, we used a rather ad hoc implementation of F_{proxy} which is only intended to evaluate the performance potential of this setting. Since the results show that it can provide performance comparable to the other two options while retaining low security requirements and engineering cost, we believe a more general solution would be desirable. Consequently, we propose Proxy+ architecture that offers fragment caching and page composition ability. Our core idea is to simply replicate the server-side caching functionality to the proxies, while carefully engineering the protocols so that consistency enforcement takes a free ride from what the programmers have already expressed when enabling server-side dynamic content caching. Our architecture requires only minor modifications to existing applications and is incrementally deployable.

We developed a prototype of Proxy+ on top of Microsoft ISA Server [15]. The proxy is augmented with a Web Filter that is able to do caching for ASP.NET applications which have made use of ASP.NET built-in output caching facility after minor modifications. In fact, these modifications will be trivial if necessary supports are absorbed by ASP.NET. Fig. 14 shows the architecture and workflow of Proxy+.

The ISA Web Filter is responsible for directing the caching of multiple versions of pages and fragments as well as composing pages. Cache is the storage of previously requested pages and fragments. Both of them reside on the ISA Server. ASP.NET Web Forms, which constitute the presentation tier of the Web application running on the server side, cooperate with the ISA Server to take advantage of its enhanced caching ability. The system workflow is described as follows:

1. An HTTP request arrives at the proxy.
2. The filter computes the cache keys of the page and its fragments.
3. If all necessary items are valid in the cache, go to step 5 and the result will be returned immediately. Otherwise, it attaches a list of keys identifying cached versions deemed relevant to the HTTP request header and forwards it to the server.

4. The application generates a (partial) response containing additional tags for delimiting cacheable fragments or behaving as placeholders to be substituted with cached content. In addition, necessary information, in the form of cache attributes tags, which allows the proxy to compute the cache keys, are sent over.
5. The filter parses the content (from the response or the cache) and fills the placeholder tags in the text with corresponding cached content, and installs any cache attributes. This way, cache attributes are incrementally pulled over on-demand.
6.
 - a. A complete response is sent back.
 - b. The fragments marked for caching are saved to the cache.

On the server side, a Web application builds its UI by using ASP.NET Web Forms classes. The content is generated by handling appropriate events in the classes. For example, “Load” is a typical event which is handled by many pages and user controls to populate content to UI.

The modification to the application must satisfy two requirements. It should be able to recognize the list of keys sent from the proxy and avoid regeneration of corresponding content. It also needs to insert additional tags to enable the proxy to do fragment caching and page assembly.

We modify the application as follows: Some new subclasses of Web Forms classes are added. All subclasses of Web Forms classes in the application will inherit from them instead. The new classes override the event dispatching and HTML outputting functions in their superclasses.

According to the key list attached in the request, they will decide whether the specific event need to be dispatched to the original handler or not (to avoid regeneration of cached content) and what additional tags and cache attributes are to be inserted into the HTML output.

The modification does not interfere with the original workflow of the application *at all*, which would greatly ease the reengineering effort. Readers are referred to [21] for more details. Although our prototype is implemented under Microsoft ASP.NET and ISA Server, we believe the method is equally applicable to other platforms, e.g., JSP.

Fig. 15 compares the average response time of Proxy+ prototype and F_{proxy} . As can be seen, when client connections are few, Proxy+ yields performance comparable to F_{proxy} , which conforms to the previous result. With more connections, the response time of Proxy+ grows faster than F_{proxy} because its feature of application independence incurs

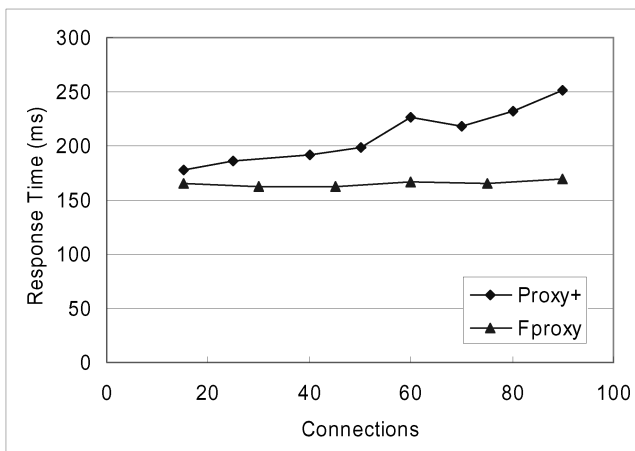


Fig. 15. Response time of Proxy+ and F_{proxy} .

more overhead in some processing on the proxy, such as parsing server responses and composing complete pages from fragments. We believe that some of these problems could be removed by a more optimized implementation.

9 SUMMARY AND CONCLUSION

We believe that Web page as a unit will disappear over time and be replaced by dynamic content that is deeply personalized. Utilizing the proxy servers located near the client to distribute and offload the processing and caching of dynamic content is entirely reasonable. This is all the more so because these edge devices are already deployed and at the same time underutilized as far as dynamic content is concerned. However, such deployment requires not only the reengineering of applications themselves, but also the considerations of security issues and careful evaluation of performance benefits.

In this paper, using a representative e-commerce benchmark, we have enumerated extensively many partitioning strategies. Without going into specifics, the following conclusions can be drawn:

1. Offloading and caching at edge proxy servers achieves significant advantages without pulling database out near the client. Our results show that, under typical user-browsing patterns and network conditions, two to three folds of latency reduction can be achieved. Furthermore, more than 70 percent of server requests are filtered at the proxies, resulting in significant server load reduction.
2. This benefit can be largely achieved by simply caching dynamic page fragments and composing the page at the proxy. Advanced offloading options can gain slightly more benefit, but are often more complex and can be counterproductive if not done carefully.
3. Many progresses in the most recent Web programming platform play important roles for this to happen. In the .NET framework, the output caching capability, stored procedures and, to a less extent, the Remoting mechanisms all made significant contributions.

4. The three-tier Web architecture gives a general guideline but is not as helpful when actual partitioning strategy is devised.

Specifically, we have evaluated the following three partitioning strategies and our findings and recommendations are as follows:

1. Replicating all application components except database to the edge provides the best average response time and the highest reduction of back-end server load. It is also the easiest to implement. However, this would be otherwise impossible if not for a highly efficient implementation of database interaction. That is to say, this option will give only disappointing offloading performance if not for the fact that database *stored procedure* has encapsulated multiple SQL queries within one request. The restriction is that since the complete business logic is pushed to the proxy, this option is suitable only for intranet applications or other situations where end-to-end security is already in place.
2. Partitioning the application components and moving some of them on the edge offer similar response time as well as server load reduction. The prerequisites are that the application is carefully partitioned and appropriate RPC mechanism is used. However, it would require considerable engineering cost for a complex Web application that was not designed to be run in a distributed fashion to begin with. Furthermore, if it is impossible to determine the sensitivity of the processing offloaded to the proxy, then this option will require end-to-end security as well.
3. We find that simply augmenting the capability of proxy today to cache dynamic fragment and compose the complete page is also very effective in terms of latency and server load reduction. The security requirement is minimal because what get cached at proxy are contents, *not* logics. This option requires change to the original application, but the process can be reasonably automated or comes for free if standard guidelines such as ESI are followed.

Our investigation thus essentially boils down to one simple recommendation: If end-to-end security is in place for a particular application, then offload all the way up to the database; otherwise, augment the proxy with page fragment caching and page composition.

Although our experiment was carried out within the .NET framework, the conclusion should be general enough to be valid on other platforms, especially for e-commerce type of applications. Many of the .NET framework and ASP.NET features are shared by other competing platforms. Furthermore, we found that network latencies dominate the performance anyway, and the .NET framework itself contributes negligible overhead in comparison.

The TTL-bound consistency enforcement as is used in this benchmark prevails in today's Web deployment, largely due to its simplicity. If more advanced, server-driven consistency maintenances [20] are to be taken in the future, we still believe our observation to hold true in general because the cache hit ratio will remain the same for all these options. What makes the difference is mainly the amount of processing being offloaded, which is orthogonal to consistency mechanism.

Based on these conclusions, we built a more robust, application-independent implementation of caching fragments and page assembly on top of the proxy architecture. It requires only simple proxy augmentation and minor changes to server-side applications to enable dynamic content caching at the edge.

Our future work includes a number of directions. We believe while e-commerce application is interesting, they will represent only a small portion in the future given that Web services will grow to cover more Web usage scenarios. Thus, we are actively seeking new applications and repeat our investigations. In addition, we believe it is important to extend this work to include underprivileged users with slow and narrow connections to the proxies. In such cases, caching at proxy is no longer sufficient: We need such functionality to move to the client side. Our preliminary experiments have indicated that this can bring about significant benefits, even though one no longer enjoys the high caching hit ratio at proxies because of sharing among different users.

ACKNOWLEDGMENTS

The authors would like to thank members of the Media Management Group for their feedback and support. Wei-Ying Ma has been instrumental in helping this work to get started. Zhigang Hua implemented Proxy+ ISA Web Filter and helped carry out the relevant experiment. The authors would also like to thank the anonymous reviewers for their helpful comments.

REFERENCES

- [1] K. Amiri, S. Park, R. Tewari, and S. Padmanabhan, "DBProxy: A Self-Managing Edge-of-Network Data Cache," IBM Research Report, RC 22419, Apr. 2002.
- [2] P. Cao, J. Zhang, and K. Beach, "Active Cache: Caching Dynamic Contents on the Web," *Proc. IFIP Int'l Conf. Distributed Systems Platforms and Open Distributed Processing (Middleware '98)*, pp. 373-388, 1998.
- [3] J. Challenger, P. Dantzig, and A. Iyengar, "A Scalable and Highly Available System for Serving Dynamic Data at Frequently Accessed Web Sites," *Proc. ACM/IEEE Supercomputing (SC '98)*, Nov. 1998.
- [4] A. Datta, K. Dutta, H. Thomas, H. VanderMeer, Suresha, and K. Ramamritham, "Proxy-Based Acceleration of Dynamically Generated Content on the World Wide Web: An Approach and Implementation," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, pp. 97-108, June 2002.
- [5] Edge Side Includes <http://www.esi.org>, year?
- [6] L. Gao, M. Dahlin, A. Nayate, J. Zheng, and A. Iyengar, "An Application Specific Data Replication for Edge Services," *Proc. 12th Int'l World Wide Web Conf. (WWW '03)*, 2003.
- [7] IBM WebSphere Application Server, <http://www-3.ibm.com/software/webservers/appserv/>, year?
- [8] IBM WebSphere Edge Server, <http://www-3.ibm.com/software/webservers/edgeserver/>, year?
- [9] A. Iyengar and J. Challenger, "Improving Web Server Performance by Caching Dynamic Data," *Proc. USENIX 1997 Symp. Internet Technologies and Systems (USTIS '97)*, Dec. 1997.
- [10] A. Labrinidis and N. Roussopoulos, "WebView Materialization," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, pp. 367-378, May 2000.
- [11] W.S. Li, W.P. Hsuing, D.V. Kalashnikov, R. Sion, O. Po, D. Agrawal, and K.S. Candan, "Issues and Evaluations of Caching Solutions for Web Application Acceleration," *Proc. 28th Int'l Conf. Very Large Data Bases (VLDB '02)*, Aug. 2002.
- [12] Microsoft .NET Pet Shop, <http://www.gotdotnet.com/team/compare/petshop.aspx>, year?

- [13] Microsoft ASP.NET Caching Features, <http://msdn.microsoft.com/library/default.asp?url=/library/en-us/cpguide/html/cponasp cachingfeatures.asp>, year?
- [14] Microsoft ASP.NET Site, <http://www.asp.net/>, year?
- [15] Microsoft ISA Server, <http://www.microsoft.com/ISAServer/>, year?
- [16] "MSDN: An Introduction to Microsoft .NET Remoting," <http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dndotnet/html/introremoting.asp>, year?
- [17] Shunra\Cloud, <http://www.shunra.com/cloud.htm>, year?
- [18] Sun Java Pet Store, http://java.sun.com/blueprints/guidelines/designing_enterprise_applications/sample_application/functionality/index.html, year?
- [19] K. Yagoub, D. Florescu, P. Valduriez, and V. Issarny, "Caching Strategies for Data-Intensive Web Sites," *Proc. Int'l Conf. Very Large Data Bases (VLDB)*, Sept. 2000.
- [20] J. Yin, L. Alvisi, M. Dahlin, and A. Iyengar, "Engineering Server-Driven Consistency for Large Scale Dynamic Web Services," *Proc. 10th Int'l World Wide Web Conf.*, May 2001.
- [21] C. Yuan, Z. Hua, and Z. Zhang, "Proxy+: Simple Proxy Augmentation for Dynamic Content Processing," *Proc. Eighth Int'l Workshop Web Content Caching and Distribution*, Sept. 2003.

Chun Yuan received the PhD degree in computer science from the University of Science and Technology of China in 2001. He has been with the Information Management and Systems Group at Microsoft Research Asia in Beijing since 2001.



Yu Chen received the bachelor and master's degrees in computer science from Peking University, majoring in software configuration management and software process management. He is an associate researcher at Microsoft Research Asia. Previously he worked on adaptive content delivery and now he takes part in some research projects in system area.



Zheng Zhang received the undergraduate degree in 1987, one year ahead of schedule. He joined HP-Labs in 1996 with a PhD degree in electrical and computer engineering from the University of Illinois at Urbana-Champaign. His research focus was high-performance computer architecture and, in particular, the memory hierarchies and high-availability support of large-scale shared memory multiprocessors. He led several teams at HP-Labs on distributed file system and P2P research. In January 2002, he joined Microsoft Research Asia as project lead and researcher, and is the research manager of system research group in MSR-Asia. His current research interest is the practice and theory of large-scale distributed system, in particular, leveraging and advancing P2P technologies. He has more than two dozen publications in leading conferences and workshops, and also served as program committee in USENIX Mobisys, IPTPS, and WWW.



► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.