

Fractal: A Mobile Code Based Framework for Dynamic Application Protocol Adaptation in Pervasive Computing *

Hanping Lufei and Weisong Shi
Wayne State University
{hlufei, weisong}@wayne.edu

Abstract

The rapid growth of heterogeneous devices and diverse networks in our daily life, makes it is very difficult, if not impossible, to build a one-size-fits-all application or protocol, which can run well in such a dynamic environment. Adaptation has been considered as a general approach to address the mismatch problem between clients and servers; however, we envision that the missing part, which is also a big challenge, is how to inject and deploy adaptation functionality into the environment. In this paper we propose a novel application level protocol adaptation framework, Fractal, which uses the mobile code technology for protocol adaptation and leverages existing content distribution networks (CDN) for protocol adaptors (mobile codes) deployment. To the best of our knowledge, Fractal is the first application level protocol adaptation framework that considers the real deployment problem using mobile code and content distribution networks. To evaluate the proposed framework, we implement an adaptive communication optimization protocol by dynamically selecting four communication protocols, including Direct sending, Gzip, Bitmap, and Vary-sized blocking. In the comparison with the static and centralized protocol adaptation approaches, evaluation shows good results on both the client side and server side. For some clients, the total communication overhead reduces 41% compared with no protocol adaptation mechanism, and 14% compared with the static protocol adaptation approach.

1. Introduction

With the development of computer and communication technologies, more and more heterogeneous devices, like desktops, laptops, PDAs, and cellular phones are connected

to the Internet using diverse networks, like Ethernet, Wi-Fi, Bluetooth, 3G/4G wireless technology. On one hand, different technologies have different characteristics. On the other hand, a heterogeneous environment makes it possible to dynamically change between different devices and network environments. For instance, a person uses a laptop with a cable modem at home, a cell phone with 3G/4G or Bluetooth on the way to the office, a desktop with Ethernet LAN in the office and a PDA with Wi-Fi in the meeting room. Diverse network connections and heterogeneous devices demand the adaptation functionality in a distributed fashion because no one-size-fits-all single function or protocol can perform well over all these networks and devices.

It is difficult, if not impossible, to build a one-size-fit-all application or protocol which can run well in the dynamic environment. Adaptation has been considered as a general approach to address the mismatch problem between clients and servers [14, 24, 36, 51]. From the perspective of adaptation locations, some of them propose the in-network adaptation, such as CANS [14], Active Names [51], Odyssey [36], and Rover [24], which focus on how to do the adaptation step by step across an overlay path. Although the functionalities are well designed, they have not considered the deployment of chosen components (drivers in CANS [14]) across multiple nodes in the path. This is an obstacle for the wide acceptance of these approaches. Other proposals try to perform the end to end adaptation, like the static content based adaptation [33, 38], which does not take the mobility of users and dynamically changing environment into consideration. From the network OSI model's point of view, some of them work in the network layer [41], which adapts the TCP/IP protocol dynamically according to the changing situations on both ends. Although the results are promising, it is not able to handle the application level protocol adaptation which makes more sense for many overlay distributed applications, e.g., streaming multicast on the Internet.

In this paper we propose Fractal, a dynamic application level protocol adaptation approach, which uses the mobile code technology for protocol adaptation and leverages existing content distribution networks (CDN) for protocol

* This work was supported in part by Michigan Life Science Corridor under grant number MEDC-459 and Wayne State University Faculty Research Grant.

adaptors (mobile codes) deployment. The idea of protocol adaptation is based on the assumption that an application protocol is composed of a series of components, also called protocol adaptors (PAD) in the Fractal framework. When a protocol needs to be adapted, the application simply needs to add or remove some PADs into or from it. Before a mobile client starts an application session with the application server, it uses the proposed interactive negotiation protocol to negotiate with the adaptation proxy deployed close to the application server. The negotiation manager inside the adaptation proxy uses the proposed adaptation path search algorithm to find one or more appropriate PADs that should be used in the following communication between the client and the application server. Metadata about these PADs will be sent to the client by the adaptation proxy. The client is then able to retrieve the PADs, which are packaged into mobile code modules, from the CDN and starts the new protocol. Although a large amount of research on mobile code and CDN has been done, few studies have combined the advantage of both of them for the protocol adaptation purpose. Specifically our contributions of this paper include:

1. *Proposing a general framework for dynamic application level protocol adaptation* —To the best of our knowledge, Fractal is the first approach on utilization of mobile code in application level protocol adaptation. With the appearance of more and more application level protocols, such as SOAP [52], LDAP [28], and Plugins, holding all the protocol implementations locally is too expensive for some network enabled mobile devices. Dynamically retrieving the necessary protocol module in an on-demand manner is applicable for mobile hosts.
2. *Dynamically adapting at the application protocol level* — Most of proposed protocol adaptation methods [3, 22, 37, 41, 46] lie in the network layer. Such systems can cope with localized changes in network conditions but cannot adapt to variations above the network layer. Moreover, their transparency hinders composability of multiple adaptations. Fractal works in the application level so it has the overall system level view to overcome this shortcoming and can maximally adapt application level protocols which have no way to be implemented in the network layer.
3. *Leveraging CDN edgeservers for protocol adaptor delivery* — CDN has already been widely deployed on the Internet to deliver Web content. Fractal extends the utilization of the content distribution network into the field of protocol adaptation. Considering PAD as a Web object, many algorithms and approaches designed for content distribution on CDN can be seamlessly transplanted to the mobile code distribution scenario. Leveraging existing CDN platforms to deliver PADs for application servers makes our approach more compatible, applicable, and extensible.
4. *Implementing an adaptive communication optimization protocol in the context of the Fractal framework* — Many communication optimization techniques are proposed in different contexts. In our previous work [30], we systematically evaluated four algorithms and found that no single algorithm out-

performed others in all cases. Different approaches have different performance in terms of different network types, document types, and device configurations. Considering these communication optimization techniques as application level protocols, we implement Fractal in a real system that dynamically chooses different communication optimization protocols and generates the application content for different client devices and network connections. Results show that using framework greatly improves both the client side and server side performance, e.g., the system capacity, client total delay, and bandwidth requirements.

The rest of the paper is organized as follows. After a brief introduction of background in Section 2, Fractal design is depicted in Section 3. A case study that builds an adaptive communication optimization protocol is presented and evaluated in Section 4. Finally, related work and conclusions are listed in Section 5 and Section 6 respectively.

2. Background

Our work is inspired by three types of previous work: mobile code [15, 21], content distribution network [1, 27], and protocol adaptation [31, 41, 45]. In this section, we explain the general background of each related research field.

2.1. Mobile Code

Mobile code [21] is defined as the data that can be executed as a program. The code can be pre-compiled for immediate execution on the recipient's processor, compiled upon receipt for subsequent execution or interpreted. The mobile code system has been used to build a distributed processing environment that is flexible in the communication abstractions it provides to applications and to enhance existing distributed applications. For the benefit of mobile code [15], a major asset provided by code mobility is that it enables service customization. The ability to request the remote execution of code helps increase application server flexibility without permanently affecting the size or complexity of the server. In Fractal we implement each protocol adaptor as a mobile code module, which is sent and executed remotely on the client side to build a new protocol allowing the client to talk with the application server.

2.2. Content Distribution Network

Content Distribution Networks (CDN) [27] is an intermediate layer of infrastructure between origin servers and clients. CDN can achieve scalable content delivery by distributing load among its edgeservers, by serving client requests from edgeservers that are close to requests, and by bypassing congested network paths. Currently CDNs are only used to deliver Web-based content. In Fractal framework, CDN is used to deliver protocol adaptor (PAD). If we

consider the PAD as a Web-based object, most of the current techniques in CDN can be leveraged to the delivery of PAD. Fractal framework extends the utilization of CDNs from traditional Web-based content to Web-based objects like mobile code and mobile agent.

2.3. Protocol Adaptation

Changing protocols to adapt link condition and network environment is not the new idea, e.g., Reno and Vegas congestion control in TCP/IP protocol [18] is a kind of adaptation. More sophisticated protocol adaptation approaches, such as STP proposed in [41], but most of them are in the network layer which makes them hard to have a general view of the whole system status. The problem of adapting to a changing network environment is further complicated because changes in network conditions are usually transparent to higher layers of the protocol stack. When higher layers, e.g., application layer, are aware of network variation, protocol adaptation can be done more adaptively and intelligently. Based on these observations, Fractal works entirely in the application layer to adapt the application protocol according to heterogeneous client environments.

3. Fractal Design – An Application Protocol Adaptation Framework

In this section we present the design of Fractal, an application protocol adaptation framework using mobile code and content distribution network edgeservers. After an overview of the Fractal framework, we in turn cover the adaptation proxy, the interactive negotiation protocol, the application protocol adaptation approach, and finally, the mobile code security mechanism.

3.1. System Overview

Fractal works entirely at the application level and has no specific requirements about underlying network topologies, connection media types, network protocols, and client hardware configurations. As an general adaptation framework, it focuses on the protocol adaptation method which uses protocol adaptors (PADs) to describe the application protocol structure and distributes the PADs to the client by CDNs for protocol the adaptation purpose. Fractal consists of five components: *application servers*, *adaptation proxies*, *CDN edgeservers*, *Protocol adaptors (PADs)*, and *client hosts* (e.g., desktop, laptop, PDA, and so on), as shown in Figure 1. The application server is the application service provider. In order to provide the functionality to heterogeneous clients in diverse environments, the application server usually communicates with clients through different application protocols. For the same application, different con-

tent (required) generated by different protocols is called *adaptive content*. For example, the content in a Web page can be transmitted or adapted using either HTTP protocol or HTTPS protocol, which is a more secure mechanism. The HTTP and HTTPS content are called *adaptive content*, as defined earlier. In Fractal, *adaptive content* can be generated either reactively or proactively. The former is suitable for the case in which content keeps changing, e.g. a stock price web site. In this scenario, memory or hard disk space requirements are small, but the price of computing the dynamic *adaptive content* maybe high. On the contrary, the latter, where *adaptive content* is precalculated in advance and saved in memory or disk consumes less CPU and has large memory or disk space requirements. The results in Section 4 show the difference between these two approaches in terms of total time.

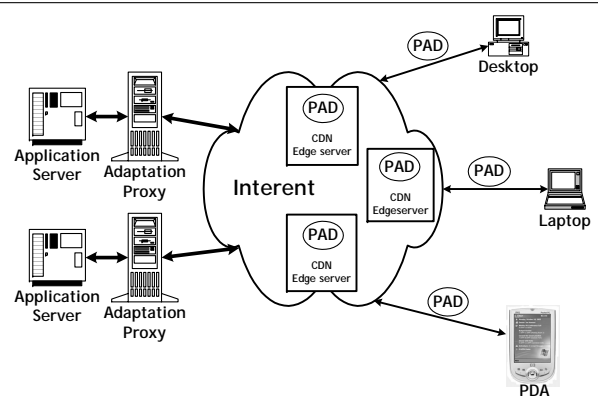


Figure 1. Architecture of application protocol adaptation using mobile code.

Although the application server can talk in many languages, i.e., protocols, the client may not have the necessary protocol to talk with the application server. To help the client talk with the application server, in Fractal we propose the notion of PAD, which is a protocol adaptor implemented in a mobile code module and deployed across the CDN edgeservers. By downloading and deploying one or more PADs, the client is then capable of starting communication with the application server using required protocols. On the server side, we assume the application server has already deployed all PADs in advance. An important issue for the client is which PADs should be used and where to find them. In the Fractal framework, close to the application server, an adaptation proxy is set up to handle the issues about PAD negotiations. Before the initialization of communication between the client and the application server, the client has to negotiate with the adaptation proxy to find proper PADs. The client will be asked to pro-

vide some metadata about his environments, such as computing ability, memory space, and network configurations to the adaptation proxy. Having these metadata, the adaptation proxy will generate the metadata of the proper PADs for the client and send the metadata of PADs back to the client. Inside these metadata is enough information for the client to download the PADs from the closest edgesever of CDNs with which the application server is associated. We will give more details about how the adaptation proxy works in the next section. Fractal leverages the wide deployment of CDNs to distribute the PADs for application servers, as illustrated in Figure 1. We envision that using CDN edgesevers for application server-specific PADs is a natural extension to the well-known Web content delivery. Note that in this paper we focus on the client/server model; however, it is straightforward to support the peer-to-peer model.

3.2. Adaptation Proxy

Adaptation proxy plays an important role in the functionalities of the Fractal framework. Usually it is deployed in the same administration domain as the application server and is responsible for negotiation with the client. A general structure of the adaptation proxy is shown in Figure 2, which includes a *negotiation manager* module and a *distribution manager* module. Each module is running as a daemon on the adaptation proxy. Next we will explain the structure and functionality of each module respectively.

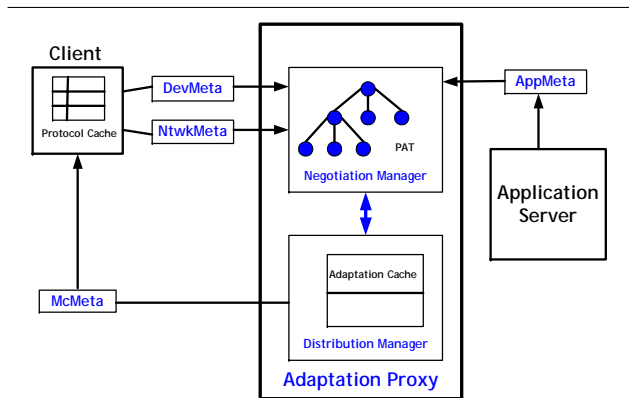


Figure 2. Structure of the adaptation proxy.

Negotiation Manager As shown in Figure 2, the negotiation manager is the key in the adaptation proxy which negotiates with the client. Some application level metadata is needed to be transmitted between the adaptation proxy and the application server, and between the adaptation proxy and the client to support the negotiation function. We define these metadata formats in Figure 3. In the rest of the

paper, we will use the acronyms in the parentheses to refer to them.

DevMeta and *NtwkMeta*, provided by clients, contain the hardware information and the network environment of the client. The application server supplies *PADMETA* to the negotiation manager, who holds the general information of each PAD. *PAD ID* is a unique identification generated by the application server. *PAD overhead* consists of the computing overhead at both the client side and server side, and corresponding traffic overhead, which is happened in the network. *Message digest* is computed using the SHA-1 [10] function and used by clients to verify the integrity of the PAD. *URL* is the link to download the PAD. Note that it is the CDN's responsibility to find the closest edgesever which holds the PAD, and to redirect the request to that edgesever. *Parent link* and *Child link* are used to build the protocol adaptation topology in the negotiation manager. *AppMeta* is comprised of *Application ID*, which marks different applications, and some *PADMETA*, which forms a protocol adaptation topology. The application server pushes new *AppMeta* to the negotiation manager when the protocol adaptation topology is first created or changed later. Usually the protocol adaptation topology is represented by a protocol adaptation tree (PAT) structure as shown in Figure 2 in the upper box located in the negotiation manager in . We will give more details about why a tree is needed and how to build and use the PAT tree in Section 3.4.1.

When the negotiation manager receives a request from a client, it first checks its adaptation cache, located in the distribution manager. The cache has entries mapping client side information to an array of *PADMETA* that the client needs. Each mapping entry is structured as follows:

$$\{ DevMeta, Application ID, NtwkMeta \} \Rightarrow \{ PADMETA_1, \dots, PADMETA_n \}$$

If the adaptation cache does not have the entry corresponding to the client side metadata, the negotiation manager then will use the algorithm described in Section 3.4.2 to form a new entry and transfer it to the distribution manager.

Distribution Manager The distribution manager is in charge of further processing of these *PADMETA* received from the negotiation manager, updating the adaptation cache, and finally sending *PADMETA* back to the client. When the distribution manager receives the *PADMETA* generated by the negotiation manager, it inserts message digest and URL data into the *PADMETA* and hides the parent and child links since the exposure to the client is unnecessary. After the negotiation procedure, which will be discussed in the following section, the distribution manager will update the adaptation cache so that the negotiation result can be directly retrieved from the cache if the same client configuration occurs later. Finally the distribution manager will handle the network communication details and send these *PADMETA* back to the client.

```

Device Metadata (DevMeta) = { Operating system type, CPU type, CPU speed, memory size }

Network Metadata (NtwkMeta) = { Network type, Network bandwidth }

PAD Metadata (PADMeta) = { PAD ID, PAD size, PAD overhead, Message digest, URL, Parent link, Child link, ... , Child link }

Application Metadata (AppMeta) = { Application ID, PADMeta 1, ... , PADMeta n }

```

Figure 3. Definitions of metadata.

Next we will explain the interactive negotiation protocol.

3.3. Interactive Negotiation Protocol

In Fractal, an interactive negotiation protocol is proposed for the interactions among these components, as shown in Figure 4. We assume both the client side and server side understand the protocol definitions. The application server has pre-deployed PADs in the application context and already pushed the *AppMeta* to the adaptation proxy, which has built a PAT inside the negotiation manager. The PADs have been distributed across the CDNs edgeservers.

At the beginning of the negotiation, a client first checks its own protocol cache, which contains some *PADMeta* saved for previous requests. If there is an entry of the protocol cache which matches the current request, the client will directly start the application communication with the application server. If not, the client sends *INIT_REQ*, which contains application request in payload, to the adaptation proxy¹ to initialize the protocol negotiation. Each packet has an *INP header* segment, which is used to maintain the interactive negotiation protocol integrity, and we will omit the details in the *INP header*. The adaptation proxy then sends *INIT_REP* as well as *ClI_META_REQ*, having empty *DevMeta* and *NtwkMeta* to be filled by the client, to acknowledge the request and ask some information about the client. After getting the reply, the client gets the content of *DevMeta* and *NtwkMeta* locally by probing the system using system calls and sends out the *ClI_META_REP*. Based on the *ClI_META_REP*, *PADMeta* is computed and sent back to the client in *PAD_META_REP* by the adaptation proxy. Next, the client updates his protocol cache and sends *PAD_DOWNLOAD_REQ* containing PAD ID to the URL of the PAD. The CDN will automatically choose a close CDN edgeserver and send back the PAD code in *PAD_DOWNLOAD_REP*. If multiple PADs are required, it is not necessary that those PADs downloaded from the same edgeserver. It is up to the CDN to manage the delivery of PADs. After the security check and PAD(s) deployment, the

client sends out the *APP_REQ* to the application server. The *APP_REQ* contains the application request as well as the negotiated protocol identifications, which notify the application server to choose the proper PADs to talk with the client. From now on the client and the application server continue the application session using the negotiated protocol. The formats of all message types used in INP are listed on the bottom of Figure 4.




Figure 4. The Interactive Negotiation Protocol.

3.4. Application Protocol Adaptation

Application protocol adaptation is the major function of Fractal. After introducing the structure and components of Fractal, we will show how the application protocol adaptation works. First, we will explain the protocol adaptation topology, the protocol adaptation tree (PAT), which is the main data structure in the procedure of adaptation. Then we will clarify the adaptation path search algorithm.

¹ Note that the client does not have to realize the existence of the adaptation proxy. The application server will automatically redirect the request to its corresponding adaptation proxy.

3.4.1. Protocol Adaptation Tree Figure 5 shows an example of the protocol adaptation tree (PAT), which is built by the negotiation manager based on *AppMeta* received from the application server. Each node of PAT is a protocol adaptor. The child PAD is an auxiliary component of the parent PAD. In order to run the parent PAD, one and only one of the children PADs must work together with the parent PAD. For example, in Figure 5, if PAD2 is the FTP protocol, PAD7 is the TCP protocol, and PAD8 is the UDP protocol, the PAD2 can choose either PAD7 or PAD8, but not both. In the real application, it is possible that one PAD is needed by multiple PADs, like TCP protocol is needed by both FTP and HTTP protocols. For the purpose of maintaining the tree structure, we use a symbolic copy of the child PAD if it is required by more than one parent PAD. For instance, in Figure 5, PAD6 is a symbolic link of PAD7, which is needed by both PAD1 and PAD2. So in order to satisfy an application protocol, a path should be found from the root application to one leaf, e.g., the path composed of PAD2 and PAD7 in the dotted line in Figure 5. Tree structure makes it flexible enough to extend adaptation protocols by adding new PAD nodes later. For example, if a new PAD, which supports PAD3, is needed later, we just add this new PAD as the first child of PAD3. Adding a new PAD in the middle, instead of the leaf of the tree, can also be done in reasonable time. From the knowledge of data structure and graph theory, we know that the number of possible paths equals the number of leaves in the tree. Next, we propose an adaptation path search algorithm to find the path.




Figure 5. The protocol adaptation tree.

3.4.2. Adaptation Path Search Algorithm The goal of the adaptation path search algorithm is to find some PADs from PAT to form an adaptation path for a client. Introducing a new protocol into an existing application will inevitably have two effects. First is the traffic overhead, which

is either increasing or reducing. Second is the extra computing overhead on both the server side and client side.

Before we choose the proper PADs for a client, the total overhead including traffic and computing overhead of each PAD is the metrics we should quantify. Running each PAD on each client configuration and each network environment to get the overhead is not a wise solution. Instead, we use a linear model and a normalized ratio to estimate these overheads. Our linear model comes from the observation that the computing overhead of each PAD is roughly proportional to the processor speed, and the traffic overhead is proportional to the network bandwidth. If the computing overhead of a PAD on one processor speed is known, the computing overhead on another processor can be deducted from the linear ratio of the speed of these two processors. Similarly we can get the traffic overhead of a PAD based on the value of another PAD and the ratio of the bandwidth of two networks. However, this linear model is not so accurate because other parameters of the processor and networks introduce error into the linear model. For example, a scientific computing module is awkward for a no floating instructor processor. A media stream application runs fluently in LAN but not in Dialup. Furthermore, an operating system is also an influential issue that we have to consider besides the linear model. For example, Microsoft DCOM can run on Windows platforms but not Unix environment. In this paper we abstract normalized ratio parameters about three key properties: *processor types*, *operating system*, and *network types* as shown in the normalized ratio matrix in the following context. Note that it is easy to introduce more parameters if necessary, e.g., the screen resolution.

As shown in Equation 1, each application server maintains the following information. $PAD_{traffic}$ is the traffic overhead of the PAD based on a standard network bandwidth, $Std_{bandwidth}$, 1Mbps, and a fixed size of traffic, 1MB in our implementation. PAD_{size} is the size vector of each PAD. PAD_{comp}^{client} is the computing overhead of PAD on a standard processor speed, Std_{cpu} , 500MHz Pentium IV in our implementation, on the client side. PAD_{comp}^{server} is the computing overhead of the PAD on the server side, which is supposed to be available in advance. All these metrics can be computed in advance. Later we will compute the estimated overhead of each PAD (PAD_{total}) for a client with specific processor speed and network bandwidth using the linear model plus the normalized ratio matrix. Specifically, we use normalized ratio matrix \mathcal{A} , \mathcal{B} , and \mathcal{R} , as shown in Equation 2, to measure the performance ratios of n number of PADs on a kinds of processor types, on b number of operation system types, and in r types of network environments. For example,

$$\begin{array}{cc}
 & \begin{array}{cc} WinCE & PalmOS \end{array} \\
 \begin{array}{c} WinMedia \\ Kinoma \end{array} & \begin{pmatrix} 1 & \infty \\ \infty & 1 \end{pmatrix}
 \end{array}$$

the above matrix shows the impacts of two operating systems (the top line) on two multimedia players (the left most column). The values in the matrix mean the Windows Media works fine in the WinCE operating system (WinCE) [55] but not in PalmOS, while Kinoma player [26] runs well in PalmOS instead of WinCE. The value of ratios does not have to be an integer. Suppose now we are about to find the better one in terms of the computing time from these two players on WinCE platform. We get the time value using the linear method as, for instance, 5 sec for WinMedia and 2 sec for Kinoma. Without the normalized matrix, Kinoma will be chosen as the better player; however, the fact is that Kinoma can not run on WinCE at all. To get the correct result, we can use the first column of this normalized matrix to adjust the linear results by multiplying 2 sec with ratio 1 for WinMedia and multiplying 5 sec with ratio ∞ for Kinoma. Then the computing time of Kinoma becomes ∞ , which immediately disqualifies itself.

$$\begin{aligned}
 \text{PAD}_{\text{traffic}} &= \begin{pmatrix} \text{pad}_1^{\text{traffic}} \\ \text{pad}_2^{\text{traffic}} \\ \vdots \\ \text{pad}_n^{\text{traffic}} \end{pmatrix}, & \text{PAD}_{\text{size}} &= \begin{pmatrix} \text{pad}_1^{\text{size}} \\ \text{pad}_2^{\text{size}} \\ \vdots \\ \text{pad}_n^{\text{size}} \end{pmatrix}, \\
 \text{PAD}_{\text{comp}}^{\text{client}} &= \begin{pmatrix} \text{pad}_1^{\text{cli-comp}} \\ \text{pad}_2^{\text{cli-comp}} \\ \vdots \\ \text{pad}_n^{\text{cli-comp}} \end{pmatrix}, & \text{PAD}_{\text{total}} &= \begin{pmatrix} \text{pad}_1^{\text{total}} \\ \text{pad}_2^{\text{total}} \\ \vdots \\ \text{pad}_n^{\text{total}} \end{pmatrix}, \\
 \text{PAD}_{\text{comp}}^{\text{server}} &= \begin{pmatrix} \text{pad}_1^{\text{svr-comp}} \\ \text{pad}_2^{\text{svr-comp}} \\ \vdots \\ \text{pad}_n^{\text{svr-comp}} \end{pmatrix}
 \end{aligned} \tag{1}$$

$$\begin{aligned}
 \mathcal{A} &= \begin{matrix} & \text{cpu}_0 & \dots & \text{cpu}_a \\ \text{pad}_0 & \begin{pmatrix} \alpha_{00} & \dots & \alpha_{0a} \\ \vdots & \ddots & \vdots \\ \alpha_{n0} & \dots & \alpha_{na} \end{pmatrix} \\ \vdots & \\ \text{pad}_n & \end{matrix}, \\
 \mathcal{B} &= \begin{matrix} & \text{os}_0 & \dots & \text{os}_b \\ \text{pad}_0 & \begin{pmatrix} \beta_{00} & \dots & \beta_{0b} \\ \vdots & \ddots & \vdots \\ \beta_{n0} & \dots & \beta_{nb} \end{pmatrix} \\ \vdots & \\ \text{pad}_n & \end{matrix}, \\
 \mathcal{R} &= \begin{matrix} & \text{ntwk}_0 & \dots & \text{ntwk}_r \\ \text{pad}_0 & \begin{pmatrix} \gamma_{00} & \dots & \gamma_{0r} \\ \vdots & \ddots & \vdots \\ \gamma_{n0} & \dots & \gamma_{nr} \end{pmatrix} \\ \vdots & \\ \text{pad}_n & \end{matrix}
 \end{aligned} \tag{2}$$

$$\begin{aligned}
 \text{PAD}_{\text{total}} &= \frac{\text{PAD}_{\text{size}}}{\text{Cli}_{\text{bandwidth}} * \rho} + \text{PAD}_{\text{comp}}^{\text{server}} \\
 &+ \frac{\text{Std}_{\text{cpu}}}{\text{Cli}_{\text{cpu}}} * \begin{pmatrix} \alpha_{0i} & 0 & \dots & 0 \\ 0 & \alpha_{1i} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \alpha_{ni} \end{pmatrix} \\
 &* \begin{pmatrix} \beta_{0j} & 0 & \dots & 0 \\ 0 & \beta_{1j} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \beta_{nj} \end{pmatrix} * \text{PAD}_{\text{comp}}^{\text{client}}
 \end{aligned}$$

$$+ \frac{\text{Std}_{\text{bandwidth}}}{\text{Cli}_{\text{bandwidth}}} * \begin{pmatrix} \gamma_{0k} & 0 & \dots & 0 \\ 0 & \gamma_{1k} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \gamma_{nk} \end{pmatrix} * \text{PAD}_{\text{traffic}}, \tag{3}$$

We abstract Equation 3 to describe the total overhead for each PAD in one client environment, $\text{PAD}_{\text{total}}$, in Equation 1. Cli_{cpu} is the device CPU speed in MHz, Cli_{mem} is the device memory size in MB, and $\text{Cli}_{\text{bandwidth}}$ is the client network bandwidth in Kbps. They come from the *DevMeta* from the client. On the right side of Equation 3, the first part is the overhead of downloading the PAD. The parameter ρ is used to capture the available application level bandwidth in a real network deployment. It is usually between 0.6 to 0.8, depending on different network types. Based on our observation, we approximate ρ as 0.8 in our design. The second part is the computing overhead on the server side. This matrix can be achieved by pre-testing each PAD on the application server. The third part is the computing overhead of running PAD on the client side. Suppose one client uses processor type i , operating system type j , and network type k , the algorithm finds the corresponding ratio vector $(\alpha_{0i} \ \alpha_{1i} \ \dots \ \alpha_{ni})^T$, $(\beta_{0j} \ \beta_{1j} \ \dots \ \beta_{nj})^T$, and $(\gamma_{0k} \ \gamma_{1k} \ \dots \ \gamma_{nk})^T$ from \mathcal{A} , \mathcal{B} , and \mathcal{R} based on its processor, operating system and network types. Given that we have only a limited number of consumer-used processors, OSes, and network types, the vector will be found with high probability. Otherwise a similar type with close parameters will be chosen instead. Then these vectors are extended to diagonal matrices, which are plugged into the Equation 3 to adjust the linear estimation. The last part of the equation is the transmission overhead of running the PAD.

After we define the approach to calculate the total overhead of the PAD, the adaptation path search algorithm starts the first step by marking each node in the PAT with the total overhead computed by Equation 3. An example is shown in Figure 5. The number beside each node is the estimated total overhead. Infinity means that the PAD is not suitable for this client environment. Then the algorithm uses the Depth-First-Search-like algorithm to traverse each path from root to leaves and finds the path with the least sum of each PAD's total overhead. The PADs on this path are the negotiated protocol result for this client. The pseudo code of the algorithm is shown in Figure 6. Take Figure 5 as an example, after line 3 in the pseudo code in Figure 6, the algorithm finishes marking each node with the total overhead shown as the number beside each node, the first path it examines is PAD1 and PAD4 and gets the *Least_Totoal_overhead* as 14 in code line 17, which is the selected shortest path so far, but when the algorithm searches along PAD2 and PAD7 with the *Least_Totoal_overhead* as 9, this new path becomes the shortest path and remains until the end of the search. Fi-

nally PAD2 and PAD7 form the final output path of the algorithm.

```

1  APISA ( IN: DevMeta, NtwkMeta, PADMeta ... PADMeta
2  OUT: PADMeta stack )
3  {
4  Use DFS to mark each node with its total overhead;
5  Least_Total_overhead = ∞;
6  path_total_overhead = 0;
7  v = root;
8  Create stack s;
9  s.push (v);
10 path_total_overhead += v.overhead;
11 mark v as visited;
12 while ( ! s.isEmpty() )
13 {
14     if (no unvisited nodes are adjacent to the node on
15     the top of the stack )
16     {
17         if ( node on top of the stack is leaf
18         && Total_overhead < Least_Total_overhead )
19         {
20             Least_Total_overhead = Total_overhead;
21             PADMeta stack = s;
22         }
23         s.pop ();
24         path_total_overhead -= node on the top.overhead;
25     }
26     else
27     {
28         select an unvisited node u adjacent to the node
29         on the top of the stack;
30         s.push (u);
31         path_total_overhead += u.overhead;
32         mark u as visited;
33     } // end if
34 } // end while
35 }

```

Figure 6. The pseudo code of the adaptation path search algorithm.

3.5. Mobile Code Security

PAD, the protocol adaptor, is the key element of the Fractal framework, and is implemented using mobile code. Security is a serious concern when deploying and running the PADs across heterogeneous environments, because the executable mobile code could possibly be written by a malicious user and allow an attacker to run native code that is subject to neither restrictions nor access control on the executing machine. In Fractal there are two techniques for securing PADs. First, sandbox [16], also known as virtual machine monitor techniques (VMM) [47], is needed to limit the privileges of PADs. The second technique used in Fractal is to assure that the source of the PAD is trustworthy using code-signing [35], in which the client manages a list of entities that it trusts. When a PAD is received, the client verifies that it was signed by an entity on this list. More advanced security techniques can be applied here, but it beyond the scope of this paper.

4. Case Study: Adaptive Communication Optimization Protocol

To evaluate the effectiveness and efficiency of the proposed Fractal framework, we implement an adaptive communication optimization protocol prototype as a case study. The basic idea of the adaptive communication optimization is to dynamically select different communication protocols, including Direct sending, Gzip, Vary-sized blocking [34], Bitmap [29], to adapt to different network conditions. This application is motivated by our recent analysis of four different communication optimization algorithms [30], in which we found that different communication optimization techniques exhibit different performance in different network environments as well as for different document types. These techniques are good examples of protocols that reduce the overall communication overhead, and inspire us to use this case to test Fractal. In the following context, we first briefly introduce each communication optimization protocol, followed by experiment platforms, the specific protocol adaptation model, and result analysis.

4.1. Four Communication Optimization Protocols

Several application-specific optimization techniques have been proposed in different contexts. Generally, they work in two fashions to reduce bandwidth requirement. One is to compress content at the server side and decompress at the client side. The other is to calculate the difference between old and new versions of the content on the server side, send difference to the client and rebuild the new version based on the difference received by the client and the old version that the client already had. In the section, we examine the four communication optimization protocols used in our case study.

1. *Direct sending* In this protocol, strictly speaking, there is no communication optimization technique, client and Web server just directly send content to each other. In this simple case, the client still needs to negotiate with the adaptation proxy at the beginning.
2. *Gzip* In this algorithm, we use gzip to compress the Web page at the Web server and decompress it at the client side. Gzip is a popular data compression program [20] which uses the LZ77 algorithm.
3. *Vary-sized blocking* Proposed in LBFS [34] for reducing traffic further, the idea of LBFS is that of content-based chunk identification. Files are divided into chunks, demarcated by points where the Rabin fingerprint [42] of the previous 48 bytes matches a specific polynomial value. This tends to identify portions even after insertions and deletions have changed its position in the file. The boundary regions are called breakpoints. The server generates the difference between two versions of a file by comparing the digest of each chunks and saves the different chunks. It is powerful to reduce the size of the difference but with expensive comput-

ing overhead on both sides. *Vary-sized blocking* has been adopted by several projects as well [7, 9, 32, 49].

4. *Bitmap* Proposed in [29], the idea behind *Bitmap* is that files are updated by dividing both files into fix-sized chunks. The client sends digests of each chunk to the server, and the server responds only with new data chunks. Based on the old version and the differencing, the new version can be rebuilt. It has outperforming results compared with other differencing algorithms for some image formats like DICOM [8], BMP, and so forth.

Basically, our evaluation results in [30] show that no single algorithm outperforms others in all cases. Different approaches have different performance in terms of different metrics. A completely different result can be achieved by the same algorithm when it is applied against different types of documents. Network bandwidth affects the performance of algorithms substantially as well. The performance can also be influenced by different parameter settings of the same algorithm. More details can be found in [30].

4.2. Experimental Platform

In our experiments platform, as shown in Figure 7, three kinds of client hosts, *desktop*, *laptop* and *Pocket PC*, use three types of network connections, *LAN*, *Wireless LAN* and *Bluetooth*, to connect to an application server and an adaptation proxy. The hardware and software configurations of the servers and clients are also shown in Figure 7. The application server holds a set of 75 Web pages with the average size of about 135KB consisting of 5KB text and four images totalling about 130KB, which is inspired by a typical example of a medical application server that holds four images of different 3D views [29]. We use Java to implement four communication optimization techniques as four protocol adaptors. The summary of function and implementation of each PAD is shown in Table 1. We also implement an adaptation proxy connected with the application server in the same LAN domain. To emulate the behavior of the real content distribution network and edgesevers, we utilize some nodes from PlanetLab [40] as the distributed PAD servers. PlanetLab has been accepted as a good platform to deploy academic-oriented CDNs platforms, such as CoDeeN [53] and Coral [13]. We set up a centralized PAD server which holds all the PADs for the purpose of performance comparisons between centralized and distributed PAD servers.

4.3. Experimental Adaptation Model

Following the Fractal framework, we first define the PADs used in this application and construct the PAT for this case study, as shown in Figure 8. The PAT in this case study is a one-level tree. Each leaf is a communication optimization PAD that can be used on a specific client envi-

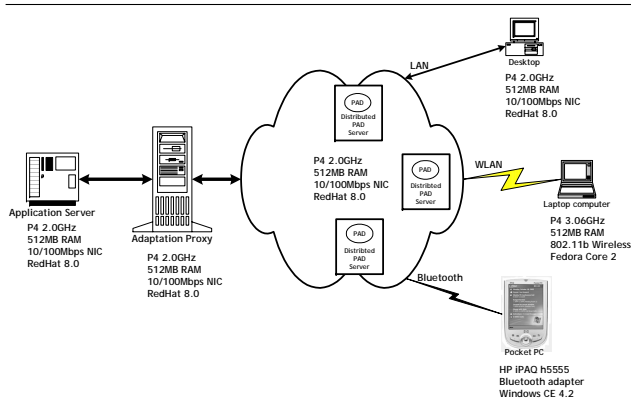


Figure 7. Configurations of experimental platform.

ronment to reduce the total communication time between the client and the application server. Then we follow Equation 3 to generate the specific Equation 7 for this case study. We use pad_{direct}^{total} , pad_{direct}^{size} , $pad_{direct}^{svr-comp}$, $pad_{direct}^{cli-comp}$, and $pad_{direct}^{traffic}$ to represent five parameters of *Direct sending* PAD: the total time overhead defined as the time from the start of downloading the PAD to the end of the application session, the size of the PAD, the server side computing overhead, the client side computing overhead, and the traffic overhead generated by the PAD. For other PADs the definitions are similar. Note that protocols like *Vary-sized blocking* and *Bitmap* have to compute the difference on the server side and rebuild a new version on the client side to reduce the bandwidth requirement. For this case study we use the normalized ratio matrix \mathcal{A} , \mathcal{B} , and \mathcal{R} in Equation 4, 5, and 6. In Equation 4, P , D , and L represent the Intel PXA 255 processor in Pocket PC, Pentium IV 2.0GHz processor in Desktop, and Pentium IV 3.06GHz processor in Laptop respectively. Some of the data come from the test, others we set as 1 to follow the linear model.




Figure 8. Protocol adaptation tree in experimental platform.

PAD name	Function	Implementation
Direct	null	null
Gzip	Compression	Java class object
Vary-sized blocking	Differencing files using Fingerprint	Java class object
Bitmap	Differencing files bit by bit	Java class object

Table 1. The functions and implementations of PADs used in the experiments.

For a newcoming client, Fractal will find its processor type, OS type, and network type, such as, i , j , and k . Then the normalized ratio matrix can be formed by collecting corresponding columns at $\mathcal{A}(i)$, $\mathcal{B}(j)$, and $\mathcal{R}(k)$. Finally with other available client side metadata, the total time overhead of each PAD for this new client can be computed using Equation 7.

$$\mathcal{A} = \begin{matrix} & & P & D & L \\ \begin{matrix} direct \\ gzip \\ vary \\ bitmap \end{matrix} & & \begin{pmatrix} 1 & 1 & 1 \\ 1.1 & 1 & 1 \\ 1.1 & 1 & 1 \\ 1.1 & 1 & 1 \end{pmatrix} \end{matrix} \quad (4)$$

$$\mathcal{B} = \begin{matrix} & & WinCE4.2 & FedoraCore2 \\ \begin{matrix} direct \\ gzip \\ vary \\ bitmap \end{matrix} & & \begin{pmatrix} 1 & 1 \\ 1 & 1 \\ 1 & 1 \\ 1 & 1 \end{pmatrix} \end{matrix} \quad (5)$$

$$\mathcal{R} = \begin{matrix} & & LAN & WLAN & Bluetooth \\ \begin{matrix} direct \\ gzip \\ vary \\ bitmap \end{matrix} & & \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix} \end{matrix} \quad (6)$$

$$\begin{pmatrix} pad_{direct}^{total} \\ pad_{gzip}^{total} \\ pad_{vary}^{total} \\ pad_{bitmap}^{total} \end{pmatrix} = \frac{1}{Cli_{bandwidth} * \alpha} * \begin{pmatrix} pad_{direct}^{size} \\ pad_{gzip}^{size} \\ pad_{vary}^{size} \\ pad_{bitmap}^{size} \end{pmatrix} + \begin{pmatrix} pad_{direct}^{svr-comp} \\ pad_{gzip}^{svr-comp} \\ pad_{vary}^{svr-comp} \\ pad_{bitmap}^{svr-comp} \end{pmatrix} + \frac{cpu}{Cli_{cpu}} * \begin{pmatrix} \alpha_{direct}(i) & & & 0 \\ & \alpha_{gzip}(i) & & \\ 0 & & \alpha_{vary}(i) & \\ & & & \alpha_{bitmap}(i) \end{pmatrix} * \begin{pmatrix} \beta_{direct}(j) & & & 0 \\ & \beta_{gzip}(j) & & \\ 0 & & \beta_{vary}(j) & \\ & & & \beta_{bitmap}(j) \end{pmatrix} * \begin{pmatrix} pad_{direct}^{cli-comp} \\ pad_{gzip}^{cli-comp} \\ pad_{vary}^{cli-comp} \\ pad_{bitmap}^{cli-comp} \end{pmatrix} + \frac{bandwidth}{Cli_{bandwidth}}$$

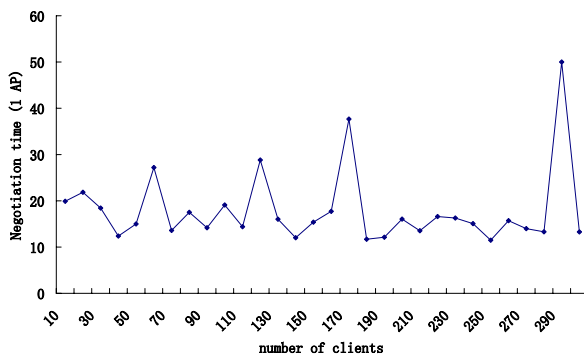
$$* \begin{pmatrix} \gamma_{direct}(k) & & & 0 \\ & \gamma_{gzip}(k) & & \\ & & \gamma_{vary}(k) & \\ & & & \gamma_{bitmap}(k) \end{pmatrix} * \begin{pmatrix} pad_{direct}^{traffic} \\ pad_{gzip}^{traffic} \\ pad_{vary}^{traffic} \\ pad_{bitmap}^{traffic} \end{pmatrix} \quad (7)$$

4.4. Results Analysis

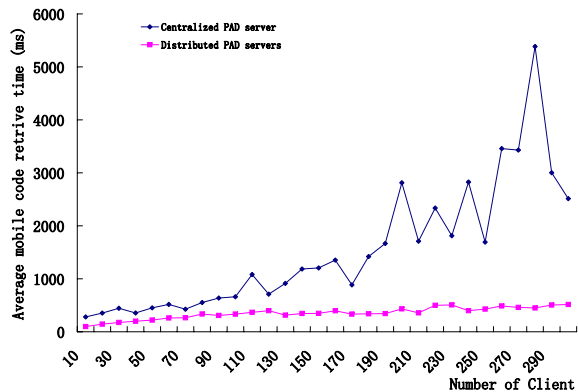
In this section, we first show the comparison of the system capacity performance between a centralized PAD server and the distributed CDN edgserver, then show the performance of protocol adaptation.

4.4.1. System Capacity Performance In Equation 3, the negotiation time, which is the time between INIT_REQ and PAD_META_REP in Figure 4 for each client, is not included because the negotiation time is not only related to PAD itself but also to the protocol adaptation topology as well as the workload of the adaptation proxy. Figure 9(a) shows the average negotiation time, y axis, versus the number of clients, x axis, up to 300 using one adaptation proxy. Although some fluctuations occur, the overall negotiation time remains in a relatively stable range for two reasons. First is the efficiency of the adaptation path search algorithm. Second is that each client only needs one time negotiation in the same environment and the application session. In order to show the benefit of deploying PADs into CDN edge-servers, we compare average PAD retrieval time in two scenarios: centralized case, in which up to 300 clients connect to the centralized PAD server simultaneously to download the PAD, and distributed case, where request traffic from same number of clients is balanced to the distributed PAD servers on the PlanetLab to simulate the CDNs and edge-servers. Figure 9(b) shows the curve of average retrieval time to the number of clients in two scenarios. We can see that the average PAD retrieval time rapidly goes up with the increasing number of clients in centralized PAD server scenario, but it steadily keeps in a small fluctuating range to give the client a roughly same retrieval time using distributed PAD servers.

4.4.2. Protocol Adaptation Performance We test each client configuration in three adaptation scenarios: No



(a) Average negotiation time



(b) Average PAD retrieve time

Figure 9. (a) Average negotiation time, (b) Average PAD retrieve time.


protocol adaptation: There is no communication optimization protocol, the client connects to the Web server and directly receives the original Web page; Fixed protocol adaptation: All clients always use one protocol, *Vary-sized blocking*, to talk with the Web server without the negotiation procedure with the adaptation proxy; Adaptive protocol adaptation: The full function of Fractal is utilized to do the protocol adaptation.

Figure 10 shows the computing overhead in three adaptation scenarios for different client configurations. The horizontal line shows three adaptation scenarios with the selected protocol in the parentheses and the vertical line, representing the computing overhead, consists of several components respectively. Figure 11(a) illustrates the bandwidth requirement in KBytes on the y axis for each client environment as shown in the x axis. We assume different clients perform identical application requests. The same protocol should generate the same number of bytes transferred, no matter the kind of client environment. First let us look at Figure 10(a), (b), and (c), which include both server side and client side computing overheads. The server side computing is used by the application server to dynamically encode the application content, e.g., compute the difference between two versions of Web pages. The client side computing overhead is used to decode the application content, e.g., rebuilding new version based on the difference and old version. *Vary-sized blocking* has huge server side computing time, which disqualifies it as the adaptive protocol for any of the client environment even if it generates the least transfer bytes as shown in Figure 11(a). Different client configurations result in different negotiated protocols, such as *Direct sending* for desktop in LAN, *Gzip* for laptop in Wireless LAN, and *Bitmap* for PDA in Bluetooth.

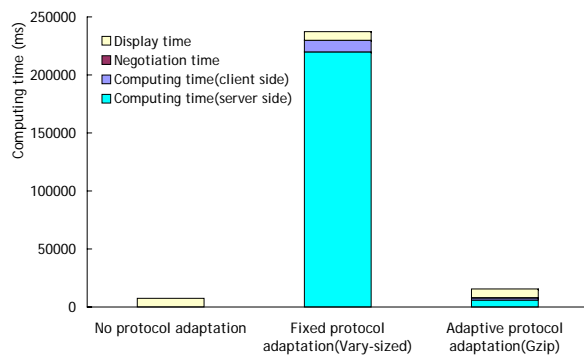
We can see that *Gzip* in Figure 10(b) and *Bitmap* in Figure 10(c) have more or less unbalanced server and client

side computing time. Since the overweighed server side computing time plays an important role in the total overhead for some protocols, e.g., *Vary-sized blocking*, different adaptation results may be observed if getting rid of the server side computing time from the total overhead. We pre-compute the server side computing tasks for each protocol on each Web page to exclude the server side computing overhead from the total computing time. We found that although the negotiated adaptation protocols for Desktop in LAN and Laptop in Wireless LAN remain the same, the adaptive protocol for PDA in Bluetooth changes from *Bitmap* to *Vary-sized blocking* as shown in Figure 10(d). Note that the scale of (c) and (d) are one order of magnitude different. The difference in negotiation results again shows that our approach can adapt the protocol according to different application strategies as well as the client environments.

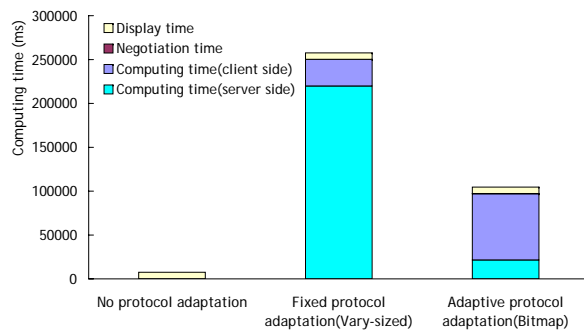
In Figure 11(a), *Direct sending* generates the most traffic bytes while *Vary-sized blocking* has the least bytes transferred. *Gzip* and *Bitmap* are in the middle in terms of bytes transferred. Computing time and bytes transferred are two components of the total overhead. In fast networks the bytes can be transferred in small time slots so that the transmission time has a smaller effect on the total overhead than the computing time. But in slow networks, bytes transferred will result into a transmission time that outweighs the computing time and dominates the total overhead. So the comprehensive influence from these two factors forms the different total overhead time performance shown in Figure 11(b) and (c). For each client configuration the adaptive protocol achieves the least total time, like *Gzip* for laptop in wireless, *Bitmap* for PDA in Figure 11(b). In the same client configuration, adaptive protocol may vary according to different server strategies, for example *Vary-sized blocking* becomes the best choice for PDA in Bluetooth without server side computing as shown in Figure 11(c). The adaptive pro-




(a) Desktop in LAN with server side computing



(b) Laptop in Wireless LAN with server side computing



(c) PDA in Bluetooth with server side computing



(d) PDA in Bluetooth without server side computing

Figure 10. A comparison of computing overhead in different environments.

ocols pointed by the oval in Figure 11 are the best choices in different scenarios, which comply exactly with the negotiation results from Fractal.

5. Related Work and Discussions

Fractal shares its goals with some recent efforts that are aimed at injecting functionality into application for adaptation. We categorize related research into four groups as *distributed adaptation*, *protocol adaptation*, *mobile code and mobile agent*, and *communication optimization*.

Distributed adaptation From the Internet topology’s point of view, adaptation functionality can be introduced either at the end-points or distributed on intermediate nodes. Odyssey [36], Rover [24] and InfoPyramid [33] are examples of systems that support end point adaptation. Conductor [56] and CANS [14] provide an application transparent adaptation framework that permits the introduction of arbitrary adaptors in the data path between applications and end services. While these approaches provide an extremely general adaptation mechanism, significant change to existing infrastructure is required for their deployment. However, Fractal solves the deployment problem by leveraging the existing CDNs technology to distributed protocol adaptors, which are implemented using mobile code.

From the network structure’s perspective, there are two issues: whether adaptation functionality is introduced at network layer with application-transparency or at the application level with application-awareness. Systems such as transformer tunnels [45] and protocol boosters [31] are examples of application-transparent adaptation efforts that work at the network level. Such systems can cope with localized changes in network conditions but cannot adapt to behaviors that differ widely from the norm. Moreover, their transparency hinders composability of multiple adaptations. More general are programmable network infrastructures, such as COMET [5], which supports flow-based adaptation, and Active Networks [48, 54], which permit special code to be executed for each packet at each visited network element. While these approaches provide an extremely general adaptation mechanism, significant change to existing infrastructure is required for their deployment. Fractal overcomes this shortcoming because it works entirely on the application level. Similar efforts also work at the application level. The cluster-based proxies in BAR-WAN/ Daedalus [11], TACC [12], and MultiSpace [17] are examples of systems where application-transparent adaptation happens in intermediate nodes (typically a small number) in the network. Active Services [2] extends these systems to a distributed setting by permitting a client application to explicitly start one or more services on its behalf that can transform the data it receives from an end service. Fractal is different from other application level frameworks in

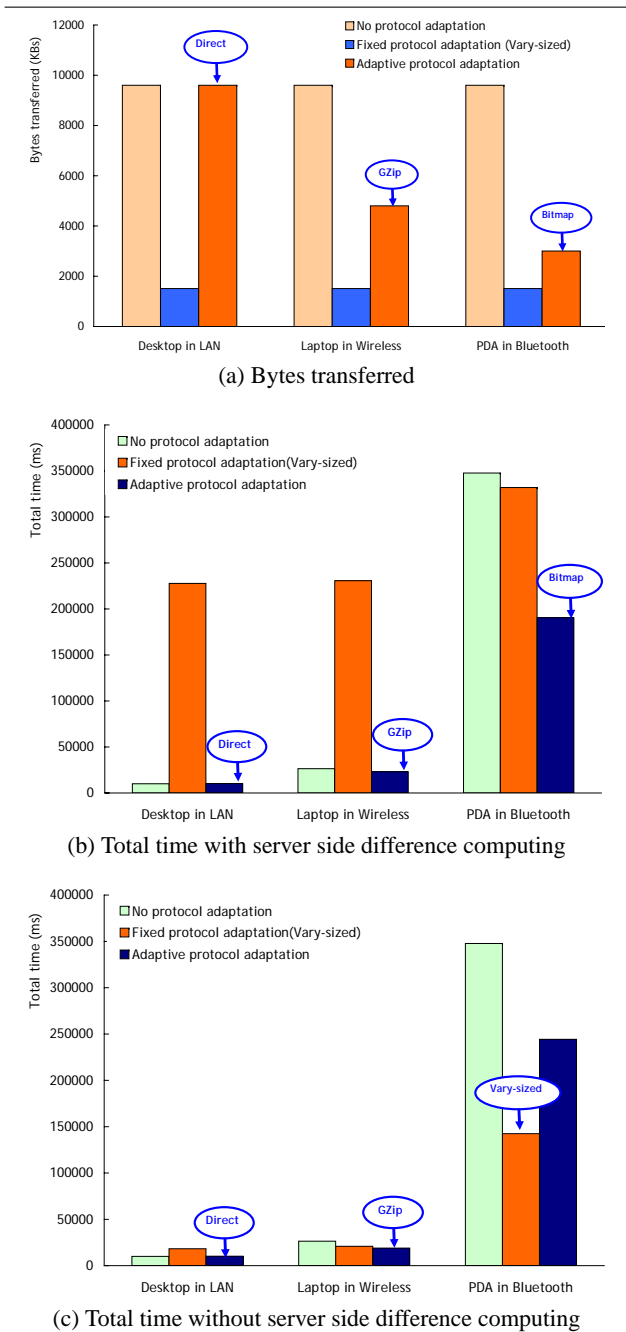


Figure 11. Comparison of three different scenarios: (a) Bytes transferred, (b) total time with server side difference computing, (c) total time without server side difference computing

the following ways: first, it is not using intermediate nodes which may occur with deployment problems. Second it does not rely on any specific data stream or client conditions. On the contrary, it is designed to cope with any applications and client environments as long as one has the proper protocol adaptor.

Protocol adaptation There are some research work about the protocol adaptation. In network level systems such as [41], in which communicating end hosts use untrusted mobile code to remotely upgrade each other with the transport protocols that they use to communicate. Transformer tunnels [45] and protocol boosters [31] are doing application-transparent adaptation by tuning the network protocol according to the change of network situations. Such systems can deal with localized changes in network conditions but cannot react to changing environments outside the network layer. Since Fractal works at the application layer, it can maximally adapt application level protocols which have no way to be completed in the network layer. Fractal is also different from the Web browser plugins, e.g., Realplay, Flash, and so on. Plugin is an application component which completes part of the functionality, incapable of doing protocol adaptation. Although today some Web sites provide multiple choices of plugins to do the similar function, they still need the client to manually select one, but maybe not the best. Fractal is a general framework to adapting the functionality by means of protocol adaptation which has transparency to the client and other characteristics, such as flexibility and extensibility, which plugins do not have.

Mobile code and mobile agent Mobile code is a good candidate for carrying a protocol module since it has long been known as a mechanism for providing a late binding of function to systems [4, 23, 25]. Mobile code and related technologies also have been proposed and studied as effective means of implementing content adaptation, protocol update, and program migration in distributed applications. In [39, 41] they propose a system in which communicating end hosts use untrusted mobile code to remotely upgrade each other with the transport protocols that are used to communicate. Our work is complimentary to their work because our proposal works in the application level. A new lightweight, component-based mobile agent system that can adapt to diverse devices and features resource saving is proposed in [6]. In this system, mobile code is brought in and associated execution states of an application dynamically after migration. NWSLite [19] provides a sophisticated predicting tools for the remote code execution offloaded from mobile client to the close server. To our best knowledge, Fractal is the first framework to use mobile code to do protocol adaptation that extends the utilization of mobile code technology.

Communication Optimization As far as the communi-

cation optimization techniques go, *Fix-sized blocking* was used in the Rsync [50] software to synchronize different versions of the same data. In this approach, files are updated by dividing both files into fix-sized chunks. The client sends digests of each chunk to the server, and the server responds only with new data chunks. *Vary-sized blocking* was proposed in LBFS [34] for further reducing traffic. Recently, several projects such as CASPER wide-area file system [49], Pond prototype [43], and Pastiche backup system [7], adopt vary-sized blocking to either improve the system performance or reduce the storage requirements. Our work compliments these efforts, and the result of this paper can be applied in their work directly. Spring and Wetherall have proposed a protocol independent technique for eliminating redundant network traffic [44]. When one end wants to send data that already exists at other end, it instead sends a token specifying where to find the data at the other end.

We believe that our work makes an initial step towards using mobile code to support the application-level protocol adaptation, in which the protocol is composed of a series of protocol adaptors. These are packaged as mobile code modules and distributed by existing CDNs. Furthermore, Fractal provides a general framework for other adaptation functionality as well by extending the PAD into other adaptation functions, e.g. content adaptation.

6. Conclusions and Future Work

In this paper, Fractal, a dynamic protocol adaptation framework, is proposed to benefit the application from choosing appropriate protocols according to dynamic client devices and network environments. To the best of our knowledge, this is the first effort on protocol adaptation by means of mobile code and CDNs edgeservers. An adaptive communication optimization protocol has been built in the context of this framework. Performance comparison with other protocol adaptation approaches shows that Fractal has lightweight system overhead, small resource footprint, and noticeable client performance improvement. Our next step includes integrating Fractal with end to end service differentiation and access control in a real pervasive computing environment, distributed computer-assisted surgery [29].

References

- [1] Akamai Technologies Inc. Edgesuite services, http://www.akamai.com/html/en/sv/edgesuite_over.html.
- [2] E. Amir, S. McCanne, and R. Katz. An Active Service Framework and its Application to Real-time Multimedia Transcoding. *Proc. of the SIGCOMM'98*, Aug. 1998.
- [3] B. Baksi, R. Krishna, N. Vaidya, and D. Pradhan. Improving performance of tcp over wireless networks. *Proceedings of the 17th ICDCS*, May 1997.
- [4] A. Birrell, G. Nelson, S. Owicki, and E. Wobber. Network objects. *Software-Practice and Experience*, pp. 25(S4):87–130, Dec 1995.
- [5] A. T. Campbell et al. A Survey of Programmable Networks. *ACM SIGCOMM Computer Communication Review*, Apr. 1999.
- [6] Y. Chow, W. Zhu, C. Wang, and F. C. Lau. The state-on-demand execution for adaptive component-based mobile agent systems. *Proc. of ICPADS*, July 2004.
- [7] L. P. Cox, C. D. Murray, and B. D. Noble. Pastiche: Making backup cheap and easy. *Proc. of the Fifth USENIX Symposium on Operating Systems Design and Implementation*, Dec. 2002.
- [8] Dicom standard, <http://medical.nema.org>.
- [9] F. Douglis and A. Iyengar. Application-specific delta-encoding via resemblance detection. *Proc. of the USENIX 2003 Annual Technical Conf.*, June 2003.
- [10] *FIPS 180-1, Secure Hash Standard*. U.S. Department of Commerce/N.I.S.T., National Technical Information Service, Springfield, VA, Apr. 1995.
- [11] A. Fox, S. Gribble, Y. Chawathe, and E. A. Brewer. Adapting to Network and Client Variation Using Infrastructural Proxies: Lessons and Perspectives. *IEEE Personal Communication*, Aug. 1998, <http://www.cs.washington.edu/homes/gribble/papers/adapt.ps.zip>.
- [12] A. Fox, S. Gribble, Y. Chawathe, E. A. Brewer, and P. Gauthier. Cluster-based Scalable Network Services. *Proc. of the 16th ACM Symp. on Operating Systems Principles*, Oct. 1997.
- [13] M. Freedman, E. Freudenthal, and D. Mazires. Democratizing content publication with coral. *Proc. of the 6th USENIX Operating Systems Design and Implementation*, Dec 2004.
- [14] X. Fu, W. Shi, A. Akkerman, and V. Karamcheti. CANS: Composable, Adaptive Network Services Infrastructure. *Proc. of the 3rd USENIX Symposium on Internet Technologies and Systems (USITS'01)*, pp. 135-146, Mar. 2001.
- [15] A. Fuggetta, G. P. Picco, and G. Vigna. Understanding code mobility. *IEEE TRANSACTIONS ON SOFTWARE ENGINEERING*, p. Vol.24 No. 5, May 1998.
- [16] L. Gong et al. Going beyond the sandbox: An overview of the new security architecture in the java development kit 1.2. *Proc. Usenix Symp. Internet Technologies and Systems, Usenix Assn.*, 1997.
- [17] S. D. Gribble, M. Welsh, E.A.Brewer, and D. Culler. The MultiSpace: An Evolutionary Platform for Infrastructural Services. *Proc. of the 1999 Usenix Annual Technical Conf.*, June 1999.
- [18] N. W. Group. Tcp slow start, congestion avoidance, fast retransmit, and fast recovery algorithms, <http://rfc.net/rfc2001.html>.
- [19] S. Gurun, C. Krintz, and R. Wolski. Nwslite: A light-weight prediction utility for mobile devices. *Proc. of MobiSys'04*, Jun 2004.
- [20] Gzip tool, <http://www.gzip.org>.

- [21] D. Halls. Applying mobile code to distributed systems, June 1997, <http://www.crema.unimi.it/mirror/scheme/thesis/node6.html>.
- [22] N. C. Hutchinson and L. L. Peterson. The x-Kernel: An Architecture for implementing Network Protocols. *IEEE Transactions on Software Engineering* 17(1):64–76, Jan. 1991.
- [23] A. D. Joseph, A. F. deLespinasse, J. Tauber, D. Gifford, and M. F. Kaashoek. Rover: a toolkit for mobile information access. *Proc. of the 15th ACM Symposium on Operating Systems Principles*, pp. 156–171, Dec 1995.
- [24] A. D. Joseph, J. A. Tauber, and M. F. Kaashoek. Mobile Computing with the Rover Toolkit. *IEEE Transaction on Computers: Special Issue on Mobile Computing* 46(3), Mar. 1997.
- [25] E. Jul, H. Levy, N. Hutchinson, and A. Black. Fine-grained mobility in the emerald system. *ACM Transactions on Computer Systems*, pp. 6(1):109–133, Feb 1988.
- [26] Kinoma player, <http://www.kinoma.com/>.
- [27] B. Krishnamurthy and J. Rexford. *Web Protocols and Practice: HTTP/1.1, Networking Protocols, Caching and Traffic Measurement*. Addison-Wesley, Inc, 2001.
- [28] Ldap (v3) revision, 2004, <http://www.ietf.org/ids.by.wg/ldapbis.html>.
- [29] H. Lufei, W. Shi, and L. Zamorano. Communication optimization for image transmission in computer-assisted surgery. *Proceedings of 2004 Congress of Neurological Surgeons Annual Meeting (abstract)*, Oct. 2004.
- [30] H. Lufei, W. Shi, and L. Zamorano. On the effects of bandwidth reduction techniques in distributed applications. *Proceedings of International Conference on Embedded and Ubiquitous Computing (EUC'04)*, Aug. 2004.
- [31] A. Mallet, J. Chung, and J. Smith. Operating System Support for Protocol Boosters. *Proc. of HIPPARCH Workshop*, June 1997.
- [32] U. Manber. Finding similar files in a large file system. *Proceedings of the USENIX Winter 1994 Technical Conference*, pp. 1-10, Jan. 1994.
- [33] R. Mohan, J. R. Simth, and C. Li. Adapting Multimedia Internet Content for Universal Access. *IEEE Transactions on Multimedia* 1(1):104–114, Mar. 1999.
- [34] A. Muthitacharoen, B. Chen, and D. Mazières. A low-bandwidth network file system. *Proc. of the 18th ACM Symp. on Operating Systems Principles (SOSP-18)*, Oct. 2001.
- [35] P. Neumann, editor. *Computer Related Risks*. Addison Wesley, 1995.
- [36] B. D. Noble. *Mobile Data Access*. Ph.D. thesis, School of Computer Science, Carnegie Mellon University, May 1998, <http://mobility.eecs.umich.edu/papers/diss.pdf>.
- [37] B. D. Noble and et.al. Agile application-aware adaptation for mobility. *Proc. of the 16th ACM Symp. on Operating Systems Principles (SOSP-16)*, Oct. 1997.
- [38] B. D. Noble, M. Price, and M. Satyanarayanan. A programming interface for application-aware adaptation in mobile computing. *Proc. 2nd USENIX Symposium on Mobile and Location-Independent Computing*, Apr 1995.
- [39] P. Patel, D. Wetherall, J. Lepreau, and A. Whitake. Tcp meets mobile code. *Proc. of the Ninth Workshop on Hot Topics in Operating Systems*, May 2003.
- [40] Planetlab, <http://planet-lab.org/>.
- [41] P. Patel, A. Whitaker, D. Wetherall, J. Lepreau, and T. Stack. Upgrading transport protocols using untrusted mobile code. *Proc. of the 19th ACM Symposium on Operating Systems Principles*, pp. 1–14, Oct 2003.
- [42] M. O. Rabin. Fingerprinting by random polynomials. Tech. Rep. TR-15-81, Harvard Aiken Computation laboratory, 1981.
- [43] S. Rhea, P. Eaton, D. Geels, H. Weatherspoon, B. Zhao, and J. Kubiatowicz. Pond: The oceanstore prototype. *Proc. of the 2nd USENIX Conf. On File and Storage Technologies*, pp. 1-14, Apr. 2003.
- [44] N. T. Spring and D. Wetherall. A protocol independent technique for eliminating redundant network traffic. *Proc. of ACM SIGCOMM'00*, pp. 87-95, Aug. 2000.
- [45] P. Sudame and B. Badrinath. Transformer Tunnels: A Framework for Providing Route-Specific Adaptations. *Proc. of the USENIX Technical Conf.*, June 1998.
- [46] P. Sudame and B. Badrinath. On providing support for protocol adaptation in mobile wireless networks. *Mobile Networks and Applications (MONET)* 6(1):43–55, 2001.
- [47] A. Tanenbaum. *Modern Operating Systems*. Second Edition, Prentice-Hall Publisher, 2001.
- [48] D. Tennenhouse and D. Wetherall. Towards an Active Network Architecture. *Computer Communications Review* 26(2), Apr. 1996, <http://www.tns.lcs.mit.edu/publications/ccr96.html>.
- [49] N. Tolia, M. Kozuch, M. Satyanarayanan, B. Karp, t. Bresoud, and a. Perrig. Opportunistic use of content addressable storage for distributed file systems. *Proc. of the USENIX 2003 Annual Technical Conf.*, June 2003.
- [50] P. Tridgell and P. Mackerras. The rsync algorithm. Tech. Rep. TR-CS-96-05, Department of Computer Science, Australian National University, 1996.
- [51] A. Vahdat, M. Dahlin, T. Anderson, and A. Aggarwal. Active names: Flexible location and transport of wide area resources. *Proc. of the 2nd USENIX Symposium on Internet Technologies and Systems (USITS'99)*, Oct. 1999, <http://www.cs.duke.edu/~vahdat/ps/an.pdf>.
- [52] W3C Consortium. Simple object access protocol (SOAP) 1.1, 2000, <http://www.w3.org/TR/SOAP/>.
- [53] L. Wang, K. Park, R. Pang, V. Pai, and L. Peterson. Reliability and security in the codeen content distribution network. *Proceedings of USENIX Annual Conference*, Jun 2004.
- [54] D. J. Wetherall, J. V. Guttag, and D. L. Tennenhouse. ANTS: A toolkit for building and dynamically deploying network protocols. *Proc. of 2nd IEEE OPENARCH*, 1998.
- [55] Windows CE Operating Systems, <http://www.microsoft.com/windowsce/>.
- [56] M. Yarvis, A. Wang, A. Rudenko, P. Reiher, and G. J. Popek. Conductor: Distributed Adaptation for complex Networks. *Proc. of the Seventh Workshop on Hot Topics in Operating Systems*, Mar. 1999, <http://lasr.cs.ucla.edu/reiher/papers/yarvis.ps>.