The Object-Oriented Design and Performance of JAWS

A High-performance Web Server Optimized for High-speed Networks

James C. Hu, Irfan Pyarali, Douglas C. Schmidt
{jxh,irfan,schmidt}@cs.wustl.edu
Department of Computer Science
Washington University, St. Louis, Missouri*

This paper appeared in the *Parallel and Distributed Computing Practices* journal, special issue on Distributed Object-Oriented Systems, edited by Maria Cobb, 1999.

Abstract

This paper provides two contributions to the study of high-performance object-oriented (OO) Web servers. First, it outlines the design principles and optimizations necessary to develop efficient and scalable Web servers and illustrates how we have applied these principles and optimizations to create JAWS. JAWS is a high-performance Web server that is designed to alleviate overheads incurred by existing Web servers on high-speed networks. In addition to its highly extensible OO design, it is also highly efficient, consistently outperforming existing Web servers, such as Apache, Java Server, PHTTPD, Zeus, and Netscape Enterprise, over 155 Mbps ATM networks on UNIX platforms.

Second, this paper describes how we have customized the JAWS OO design to leverage advanced features of Windows NT on multi-processor platforms linked by high-speed ATM networks. The Windows NT features used in JAWS include asynchronous mechanisms for connection establishment and data transfer. Our previous benchmarking studies demonstrate that once the overhead of disk I/O is reduced to a negligible constant factor (e.g., via memory caches), the primary determinants of Web server performance are its concurrency and event dispatching strategies.

Our performance results over a 155 Mbps ATM network indicate that certain Windows NT asynchronous I/O mechanisms (i.e., `TransmitFile`*) provide superior performance for large file transfers compared with conventional synchronous multi-threaded servers. Conversely, synchronous event dispatching performed better for files less than 50 Kbytes. Thus, to provide optimal performance, a Web server design should be adaptive, i.e., choosing to use different mechanisms (such as* `TransmitFile`*) to handle requests for large files, while using alternative I/O mechanisms (such as synchronous event dispatching) on requests for small files.*

1 Introduction

The emergence of the World Wide Web (Web) as a mainstream development platform has yielded many hard problems for softwar developers, who must provide high quality of service to application users. Strategies for improving client performance include client-side caching and caching proxy servers [29]. However, performance bottlenecks persist on the server-side due to factors such as inappropriate choice of concurrency and dispatching strategies, excessive filesystem access, and unnecessary data copying.

As high-speed networks (such as ATM) and high-performance I/O subsystems (such as RAID) become ubiquitous, the bottlenecks of existing Web servers become increasingly problematic. To alleviate these bottlenecks, Web servers must utilize an integrated approach that combines optimizations at multiple levels. Figure 1 illustrates the general architecture of such a Web system.

This diagram provides a layered view of the architectural components required for an *HTTP client* to retrieve an HTML file from an *HTTP server*. Through *GUI* interactions, the client application user instructs the HTTP client to retrieve a file. The *requester* is the active component of the client that communicates over the *network*. It issues a request for the file to the server with the appropriate syntax of the *transfer protocol*, in this case HTTP. Incoming requests to the *HTTP server* are received by the *dispatcher*, which is the request demultiplexing engine of the server. It is responsible for creating new threads or processes (for concurrent Web servers) or managing descriptor sets (for single-threaded concurrent servers). Each request is processed by a *handler*, which goes through a *lifecycle* of parsing the request, logging the request, fetching file

---

[*] This work was supported in part by grants from OTI, Kodak, and Siemens.

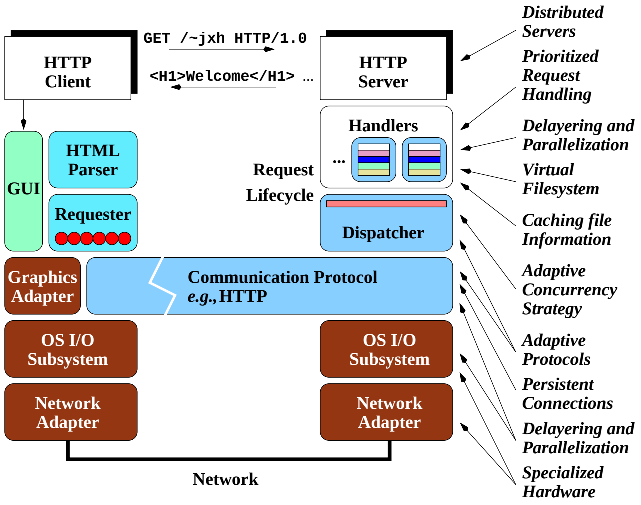

Figure 1: Overview of a Typical Web System and Optimizations

status information, updating the cache, sending the file, and cleaning up after the request is done. When the response returns to the client with the requested file, it is parsed by an *HTML parser* so that the file can be rendered. At this stage, the *requester* may issue other requests on behalf of the client, *e.g.*, in order to fill a client-side cache.

Our experience developing Web servers for multiple OS platforms indicates that the effort required to improve performance can be simplified significantly by leveraging OS features explicitly. For example, an optimized file I/O system that automatically caches open files in main memory helps to reduce latency. Likewise, support for asynchronous event dispatching [9] and the Proactor pattern [8] can increase server throughput by reducing the context switching and synchronization overhead incurred from multi-threading.

This paper presents two complementary strategies for developing optimized Web servers. First, we present empirical results demonstrating that to achieve optimal performance, Web servers must support *dynamic* adaptivity (*i.e.*, the ability to update behavior "online" to account for changes in run-time conditions). Second, we describe our recent efforts at adapting a high-performance Web server (developed originally on UNIX) to leverage the asynchronous event dispatching mechanisms on Windows NT. This work illustrates the importance of an extensible Web server design that supports *static* adaptivity, *i.e.*, changing the behavior of the Web server "off-line" to account for OS platform characteristics.

Our research vehicle for demonstrating the effectiveness of dynamic and static adaptation is *JAWS*. JAWS is both an adaptive Web server and an OO development framework for Web servers that run on multiple OS platforms, including Win32 (*i.e.*, Windows NT and Windows '95), most versions of UNIX (*e.g.*, SunOS 4.x and 5.x, SGI IRIX, HP-UX, DEC UNIX, AIX, Linux, and SCO), and MVS OpenEdition.

The need for dynamic and static adaptivity in JAWS can be motivated as follows:

The need for dynamic adaptivity: On many OS platforms, under different workloads, a single, statically configured content transfer mechanism cannot provide optimal performance. Results in this paper show that the performance of different OS level I/O mechanisms varies considerably according to file size. For instance, on Windows NT 4.0, synchronous I/O provides the best performance for transferring small files, whereas the `TransmitFile` operation provides the best performance for transferring large files under heavy loads.

The need for static adaptivity: To achieve high performance, Web servers must be adapted statically to use native high-performance mechanisms provided by the OS platform. For example, different OS platforms may provide specialized I/O mechanisms, such as asynchronous I/O or bulk data transfer, or specialized devices, such as high-speed ATM network interfaces [3]. Therefore, simply porting a Web server to use common OS mechanisms and APIs, such as BSD sockets, `select`, and POSIX threads, is not sufficient to achieve maximal performance on different OS platforms.

The results in this paper are based on extensions to JAWS' original synchronous event dispatching model, which was based on the POSIX threading model and BSD sockets. These extensions support the asynchronous event dispatching and communication mechanisms available on Windows NT (JAWS-NT). The Windows NT mechanisms incorporated into JAWS-NT include overlapped I/O, I/O completion ports, `TransmitFile`, `GetQueueCompletionStatus`, and `AcceptEx`.[1] It was fairly straightforward to customize JAWS to support the new asynchronous mechanisms because JAWS was developed as an extensible OO Web server framework.

As shown in Section 4, the performance measurements of JAWS-NT over a ∼155 Mbps ATM link indicate significant throughput and latency variance between the synchronous and asynchronous event dispatching and concurrency models on Windows NT. In addition, our experience with the Windows NT asynchronous event dispatching mechanisms has revealed other benefits besides improved throughput and latency. For instance, asynchronous event dispatching allows Web servers to significantly reduce the number of threading resources required to handle client requests concurrently.

The remainder of this paper is organized as follows: Section 2 provides an overview of the JAWS' OO server framework design and explains the optimizations we have applied to it; Section 3 outlines the concurrency strategies supported by

---

[1] These Windows NT mechanisms are described in Section 3.3.2.

JAWS-NT and describes the key differences between the synchronous and asynchronous event dispatching models; Section 4 analyzes our performance measurements of JAWS-NT over an ATM network; Section 5 compares a highly optimized JAWS implementation against Netscape Enterprise and Microsoft Internet Information Server (IIS); Section 6 compares JAWS with related work; and Section 7 presents concluding remarks.

2 The Object-Oriented Design of JAWS

The UNIX version of JAWS (described in [6]) consistently outperforms other servers in our test suite of Web servers over 155 Mbps ATM networks. This section briefly outlines the design principles and optimizations used by JAWS to achieve such high performance.

2.1 Determinants of Web Server Performance

JAWS is both a Web server and an OO framework [26] written in C++ that facilitates the development of flexible and adaptive high-performance Web systems. The optimizations, OO design principles, and patterns used in JAWS are guided by results from our empirical analysis [6, 7, 8] of Web server performance bottlenecks over high-speed ATM networks. Assuming sufficiently high network bandwidth and large file system caching, our experiments have identified the following determinants of Web server performance:

Concurrency strategy and event dispatching: Request dispatching occupies a large portion (*i.e.*, ~50%) of non-I/O related Web server overhead. Therefore, the choice of concurrency strategy, such as thread/process pool vs. thread/process-per-request, and dispatching strategy, such as asynchronous vs. synchronous, has a major impact on performance.

Avoiding the filesystem: Web servers that implement sophisticated file data and file stats caching strategies, such as PHTTPD and JAWS, perform much better than those that do not, such as earlier versions of Apache [6].

2.2 Applying Patterns and Frameworks

Developers of Web servers strive to build fast, scalable, and configurable systems. However, there are some common pitfalls encountered by these developers. Common pitfalls include (1) coping with tedious and error-prone low-level programming details, (2) lack of portability, and (3) the complexity of navigating the wide range of server design alternatives. By carefully utilizing patterns and frameworks, these hazards can be avoided, by allowing developers to leverage reuse of design and code.

Figure 2 illustrates the major structural components and design patterns that comprise the JAWS Adaptive Web Server (JAWS) framework. JAWS is designed to allow the customization of various Web server strategies in response to environmental factors. These factors include *static* factors, such as the number of available CPUs, support for kernel-level threads, and availability of asynchronous I/O in the OS and *dynamic* factors, such as Web traffic patterns and workload characteristics.

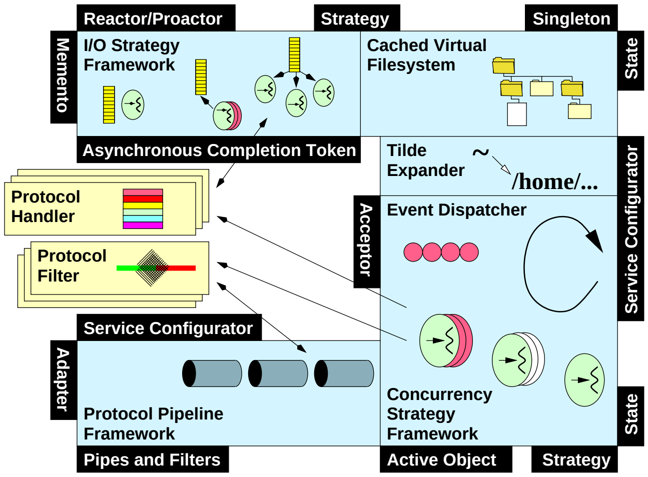

Figure 2: Architectural Overview of the JAWS Framework

2.3 Components and Patterns in JAWS

JAWS is structured as a *framework of frameworks*. The overall JAWS framework contains the following components and frameworks based on the referenced patterns:

- the *Event Dispatcher*, predominantly follows the *Acceptor* pattern,
- a *Concurrency Strategy*, using *State* and *Active Object* patterns,
- an *I/O Strategy*, incorporating *Reactor* and *Proactor* patterns,
- a *Protocol Pipeline*, which implements the *Pipes and Filters* pattern,
- the *Protocol Handlers*, which adopt the *Adapter* pattern, and
- *Cached Virtual Filesystem*, which uses the *Strategy* pattern.

Each framework is structured as a set of collaborating objects implemented using components in ACE [23]. The collaborations among JAWS components and frameworks are guided by a family of patterns, which are listed along the borders in Figure 2. An outline of the key frameworks, components, and patterns in JAWS is presented below, along with illustrations of how the patterns are implemented withing JAWS. These illustrations should be regarded as example implementations of these patterns by JAWS.[2]

Event Dispatcher: This component is responsible for coordinating JAWS' *Concurrency Strategy* with its *I/O Strategy*. As illustrated in Figure 3, the passive establishment of connection events with Web clients follows the *Acceptor* pattern [25]. New incoming HTTP request events are serviced by a concurrency strategy. As events are processed, they are dispatched to the *Protocol Handler*, which is parameterized by an I/O strategy. JAWS ability to dynamically bind to a particular concurrency strategy and I/O strategy from a range of alternatives follows the *Strategy* pattern [4].

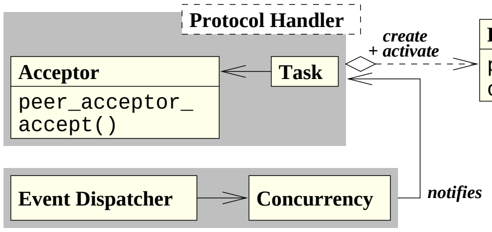

Figure 3: Structure of the Acceptor Pattern in JAWS

Concurrency Strategy: This framework implements concurrency mechanisms, such as single-threaded, thread-per-request, or thread pool, that can be selected adaptively at run-time using the *State* pattern [4] or pre-determined at initialization-time. The *Service Configurator* pattern [10] is used to configure a particular concurrency strategy into a Web server at run-time. When concurrency involves multiple threads, the strategy creates protocol handlers that follow the *Active Object* pattern [12], as illustrated in Figure 4.

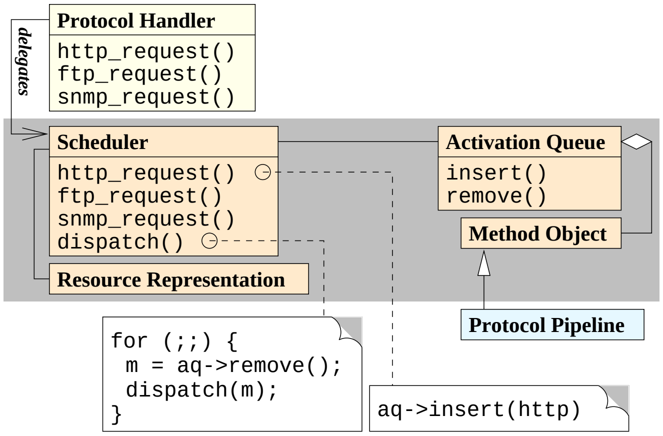

```
for (;;) {
 m = aq->remove();
 dispatch(m);
}
```

```
aq->insert(http)
```

Figure 4: Structure of the Active Object Pattern in JAWS

I/O Strategy: This framework implements various I/O mechanisms, such as asynchronous, synchronous and reactive I/O. Multiple I/O mechanisms can be used simultaneously. In JAWS, asynchronous I/O is implemented using the *Asynchronous Completion Token* [19] pattern and *Proactor* [9] pattern, as illustrated in Figure 5. Reactive I/O is accomplished through the *Reactor* pattern [24]. Reactive I/O utilizes the *Memento* pattern [4] to capture and externalize the state of a request so that it can be restored at a later time.

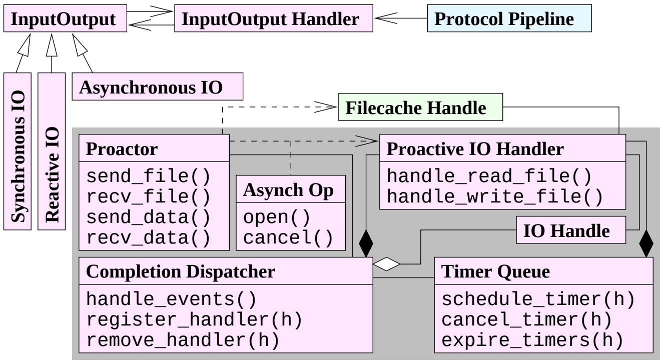

Figure 5: Structure of the Proactor Pattern in JAWS

Protocol Handler: This framework allows system developers to apply the JAWS framework to a variety of Web system applications. A *Protocol Handler* is parameterized by a concurrency strategy and an I/O strategy. These strategies are decoupled from the protocol handler using the *Adapter* [4] pattern. In JAWS, this component implements the parsing and handling of HTTP/1.0 request methods. The abstraction allows for other protocols, such as HTTP/1.1, DICOM, and SFP [18], to be incorporated easily into JAWS. To add a new protocol, developers simply write a new *Protocol Handler* implementation, which is then configured into the JAWS framework.

Protocol Pipeline: This framework allows filter operations to be incorporated easily with the data being processed by the *Protocol Handler*. This integration is achieved by employing the Adapter pattern. Pipelines follow the *Pipes and Filters*

---

[2] Due to space limitations it is not possible to describe each pattern in detail. The references provide additional information on each pattern mentioned in this section.

pattern [1] for input processing. Pipeline components can be linked dynamically at run-time using the *Service Configurator* pattern, as shown in Figure 6.

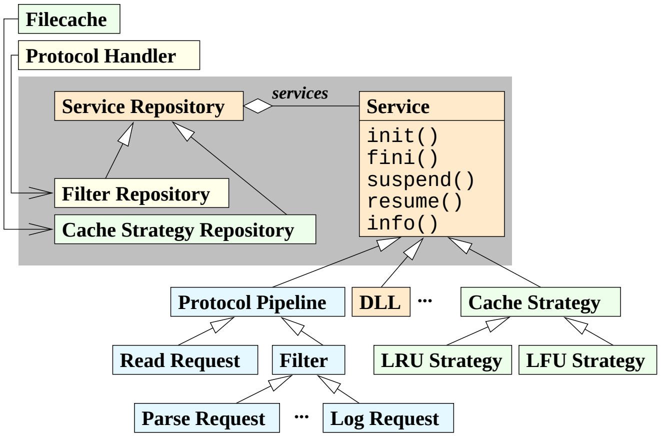

Figure 6: The Service Configurator Pattern in JAWS

Cached Virtual Filesystem: This component improves Web server performance by reducing the overhead of filesystem access. Various caching strategies, such as LRU, LFU, Hinted, and Structured, can be selected following the *Strategy* pattern [4]. This allows different caching strategies to be profiled and selected based on their performance. Moreover, optimal strategies to be configured statically or dynamically using the *Service Configurator* pattern, as shown in Figure 6. The cache for each Web server is instantiated using the *Singleton* pattern [4].

Tilde Expander: This component is another cache component that uses a perfect hash table [22] to map abbreviated user login names (*e.g.*, `~schmidt`) to user home directories (*e.g.*, `/home/cs/faculty/schmidt`). When personal Web pages are stored in user home directories, and user directories do not reside in one common root, this component substantially reduces the disk I/O overhead required to access a system user information file, such as `/etc/passwd`. By virtue of the *Service Configurator* pattern, the Tilde Expander can be unlinked and relinked dynamically into the server when a new user is added to the system.

3 Event Dispatching and Concurrency Strategies for Web Servers

The JAWS `Event Dispatcher` is a flexible component that can be configured to use multiple Concurrency Strategies, such as thread pool and thread-per-request. The initial design of JAWS used a *synchronous* event dispatching model because it was developed on Solaris 2.5, which does not provide efficient asynchronous I/O support. This section describes how the JAWS framework was enhanced to support the *asynchronous* event dispatching model provided by Windows NT 4.0.

3.1 Event Dispatching Strategies

3.1.1 Synchronous Event Dispatching

A common Web server architecture uses synchronous event dispatching. This architecture consists of two layers: the *I/O Subsystem*, and the *Protocol Handlers*, as shown in Figure 7. The I/O Subsystem typically resides in the kernel and is implemented with sockets. One socket plays the role of the acceptor, which is a factory that creates new data sockets. Protocol Handlers, having its own thread of control, reads and processes the data coming from the socket that was created from a newly accepted connection. Synchronous event dispatching dedicates the selected thread to the new client for the duration of the file transfer.

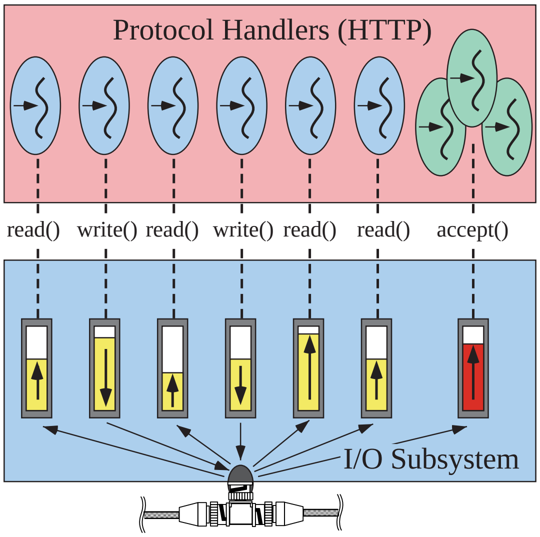

Figure 7: Synchronous Event Dispatching

3.1.2 Asynchronous Event Dispatching

The asynchronous event dispatching architecture also consists of both I/O subsystem and protocol handler layers, as shown in Figure 8. However, in asynchronous I/O, each I/O operation is “handed off” to the kernel, where it runs to completion. Thus,

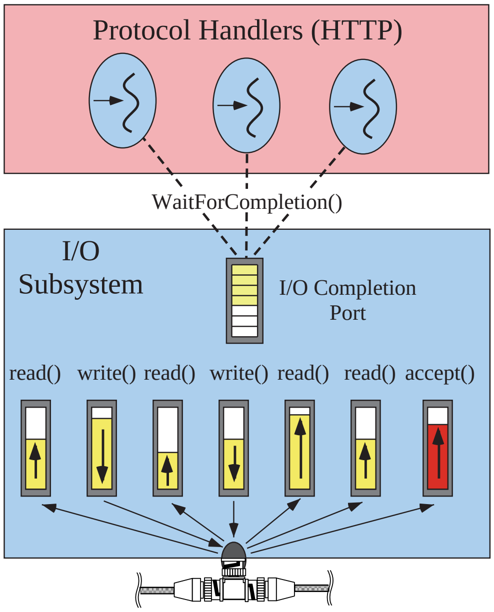

Figure 8: Asynchronous Event Dispatching

the initiating thread does not block. When the kernel has completed the operation, the kernel notifies the process through an *I/O completion port*. An I/O completion port is a kernel-level thread-safe queue of I/O completion notifications.

The primary benefits of using I/O completion ports include the following:

Increased flexibility and scalability: The thread initiating the asynchronous I/O operation and the thread dequeueing the completion status from the port can be different. This makes it possible to tune the level of concurrency in an application by simply increasing the number of completion handler threads.

Fewer threads and less overhead: The asynchronous thread pool requires significantly fewer threads than the synchronous thread pool because threads no longer block on I/O operations. The reduction of threads in the system reduces context switching and synchronization overhead.

The primary drawback of I/O completion ports is the complexity of the asynchronous programming model. Servers programmed to use I/O completion ports directly require extra data structures in addition to the run-time stack. These data structures are used to save and restore state explicitly when event completions are dispatched asynchronously.

The complexity of asynchronous I/O and I/O completion ports can be alleviated by applying the Proactor pattern [8]. This pattern supports both efficient and flexible asynchronous event dispatching strategies for high-performance concurrent applications. In general, applying this pattern enables developers to leverage the performance benefits of executing operations concurrently, without exposing the complexity of I/O completion ports and asynchronous I/O directly. In our experience, applying the Proactor pattern to the JAWS Web OO server framework made it considerably easier to design, develop, test, and maintain.

3.2 Concurrency Strategies

Existing Web servers use a wide range of concurrency strategies. These strategies include single-threaded concurrency (*e.g.*, Roxen), process concurrency (*e.g.*, Apache), and thread concurrency (*e.g.*, PHTTPD and Zeus). Single-threaded servers cannot take advantage of multi-CPU hardware concurrency. Likewise, servers that use multiple processes incur higher process creation overhead. Our discussion focuses on threading strategies because process creation incurs significantly greater overhead than thread creation.

For instance, our measurements revealed that the time required to create a thread on a 180 MHz, dual-CPU Pentium PRO2 running Windows NT 4.0 is ~0.4ms. However, the time required to create a process is ~4.5ms, which is an order of magnitude higher. As a result, JAWS implements concurrency via the threading strategies described below.

3.2.1 Thread-per-Request

In the thread-per-request model, a new thread is spawned to handle each incoming request. As shown in Figure 9, one thread blocks on the acceptor socket. This acceptor thread is a factory that creates a new handler thread to interact with each client.

After creating a new handler thread, the acceptor thread continues to wait for new connections on the acceptor socket. In contrast, the handler thread reads the HTTP request, services it, and transmits the result to the client. The lifecycle of a handler thread completes after the data transfer operation is finished.

The thread-per-request model is useful for long-duration requests from multiple clients. It is less useful for short-duration requests due to the overhead of creating a new thread for each request. In addition, it can consume a large number of OS resources if many clients simultaneously perform requests during periods of peak load.

3.2.2 Thread Pool

In the thread pool model, a group of threads are pre-spawned during Web server initialization, as shown in Figure 10. Prespawning eliminates the overhead of creating a new thread for

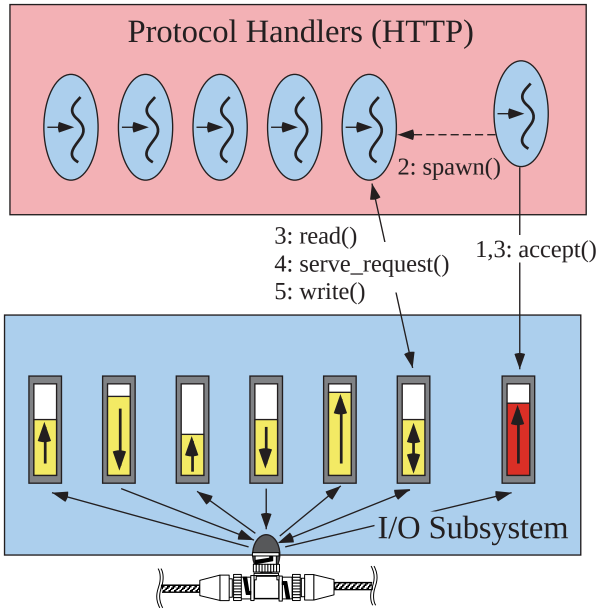

Figure 9: Thread-per-Request

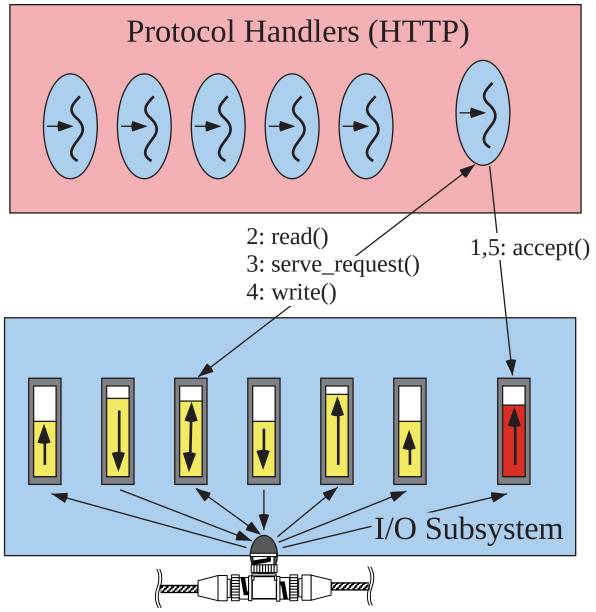

Figure 10: Thread Pool

each request. Each thread blocks in `accept` waiting for connection requests to arrive from clients. When a new connection arrives, the OS selects a thread from the pool to `accept` it and return a data socket handle.

The thread then performs a synchronous `read` from the newly connected data socket handle. Once the entire HTTP request has been read, the thread performs the necessary computation and filesystem operations to service the request. The requested data is then transmitted synchronously to the client. The thread returns to the thread pool and reinvokes `accept` after the data transmission completes.

Synchronous thread pool is useful for bounding the number of OS resources consumed by a Web server. Client requests can execute concurrently until the number of simultaneous requests exceeds the number of threads in the pool. At this point, additional requests must be queued until a thread becomes available. This queue is typically maintained in the OS kernel's TCP layer.

To reduce latency, the thread pool can be configured to always have threads available to service new requests. The number of threads needed to support this policy can be very high during peak loads because threads block in long-duration synchronous I/O operations. The asynchronous thread pool approach (described in Section 3.2.3) improves this model by considerably reducing the number of threads in the system.

Incidentally, in some operating systems, such as versions of Solaris before 2.6, it is not possible to have multiple threads in a process all blocking simultaneously in `accept` on the same acceptor socket. Therefore, on these OS platforms, it is necessary to protect `accept` with a mutex lock in a thread pool Web server. In contrast, Windows NT 4.0 supports simultaneous `accept` calls, so additional synchronization is not required in JAWS-NT.

3.2.3 Asynchronous Thread Pool

Figure 8 shows the asynchronous thread pool model. Like the synchronous model, the asynchronous thread pool is created during Web server initialization. Unlike the synchronous model, however, the threads wait on a *completion port* rather than waiting on `accept`. The OS queues up results on the completion port from all asynchronous operations, such as asynchronous accepts, reads, and writes.

The result of each asynchronous operation is handled by a thread the OS selects from the pool of threads waiting on the completion port. This asynchronous model is useful because the same programming model works for a single thread, as well as multiple threads. The thread that initiated the asynchronous operation need not be the one selected to handle its completion, however. Therefore, it is hard to implement concurrency strategies other than thread pool with an asynchronous event dispatching model.

3.2.4 Thread-per-Connection

In the thread-per-connection model the newly created handler thread is responsible for the lifetime of the client connection, rather than just a single request from the client. Therefore, the new thread may serve multiple requests before terminating. Thread-per-connection is not suitable for HTTP 1.0 because it establishes a new connection for each request. This concurrency model *is* applicable, however, in HTTP 1.1, which supports persistent connections [15, 20].

Thread-per-connection provides good support for prioritization of client requests. For instance, higher priority clients can be associated with higher priority threads. Thus, request from high priority clients will be served ahead of other requests because the OS can preempt lower priority threads.

One drawback to thread-per-connection is that if certain connections receive considerably more requests than others, they can become a performance bottleneck. In contrast, thread pool and thread-per-request provide better support for load balancing.

3.3 Optimizing JAWS for Windows NT

The performance results presented in Section 4 were conducted using a version of JAWS that was customized for Windows NT (JAWS-NT). The NT-specific optimizations applied in JAWS are described below.

3.3.1 Overview of Windows NT Asynchronous I/O

The Windows NT asynchronous I/O model supports *proactive* semantics, which allow applications to actively initiate I/O-related operations, such as `ReadFile`, `WriteFile`, and `TransmitFile`. The following steps are required to program asynchronous I/O on Windows NT:

1. Create I/O and event handles: First, a `HANDLE` is created that corresponds to an I/O channel for the type of networking mechanism, such as a socket or named pipe, used by the application.

Next, an application creates a `HANDLE` to a Win32 event object and uses this event object's `HANDLE` to initialize an overlapped I/O structure. The event object will be signaled when asynchronous operations on the `HANDLE` complete.

2. Asynchronous operation invocation: The `HANDLE` to the I/O channel and the overlapped I/O structure are then passed to the asynchronous I/O operation, such as `WriteFile`, `ReadFile`, or `AcceptEx`. The initiated operation proceeds asynchronously and does not block the caller.

3. Asynchronous operation completion: When an asynchronous operation completes, the event object specified inside the overlapped I/O structure is set to the "signaled" state. Subsequently, Win32 demultiplexing functions, such as `WaitForSingleObject` or `WaitForMultipleObjects`, can be used to detect the signaled state of the Win32 event object. These functions indicate when an outstanding asynchronous operation has completed.

3.3.2 Overview of Windows NT Functions Relevant to Web Servers

The following Win32 functions are particularly relevant for developers of asynchronous Web Servers on Windows NT:

GetQueueCompletionStatus: The function `GetQueueCompletionStatus` attempts to dequeue an I/O completion result from a specified completion port. If there are no completion results queued on the port, the function blocks the calling thread waiting for asynchronous operations associated with the completion port to finish. The blocking thread returns from the `GetQueuedCompletionStatus` function when (1) it can dequeue a completion packet or (2) when the function times out.

Windows NT selects the thread that has executed most recently from among the waiter threads on the completion port to handle the new connection. This reduces context switching overhead because it increases the likelihood that thread context information is still cached in the CPU and OS.

AcceptEx: The `AcceptEx` function combines several socket functions into a single API/kernel transition. When it completes successfully, `AcceptEx` performs the following three tasks: (1) a new connection is accepted, (2) both the local and remote addresses for the connection are returned, and (3) the first block of data sent by the remote client is received. Microsoft [14] claims that programs establishing connections with `AcceptEx` will perform better than those using the `accept` function. However, our results in Section 4 show that over high-speed networks there is not much difference in performance.

TransmitFile: `TransmitFile` is a custom Win32 function that sends file data over a network connection, either synchronously or asynchronously. The function uses the Windows NT virtual memory cache manager to retrieve the file data. As shown in Section 4, the asynchronous form of `TransmitFile` is the most efficient mechanism for transferring large amounts of data over sockets on Windows NT.

In addition to transmitting files, `TransmitFile` allows data to be prepended and appended before and after the file data, respectively. This is particularly well-suited for Web servers because they typically send `HTTP` header and trailer data with the requested file. Hence, all the data to the client can be sent in a single system call, which minimizes mode switching overhead.

The Windows NT *Server* optimizes `TransmitFile` for high performance; all our benchmarks were run on NT Server. The Windows NT *Workstation* optimizes the function for minimum memory and resource utilization. Our measurements confirm that that `TransmitFile` on Windows NT *Server* substantially outperforms `TransmitFile` on Windows NT *Workstation*.

3.3.3 Customizing JAWS for Windows NT Asynchronous I/O

We modified the original Solaris implementation of JAWS to use the Windows NT asynchronous event dispatching model described above. Figure 11 shows the interactions between components in the resulting JAWS-NT model.

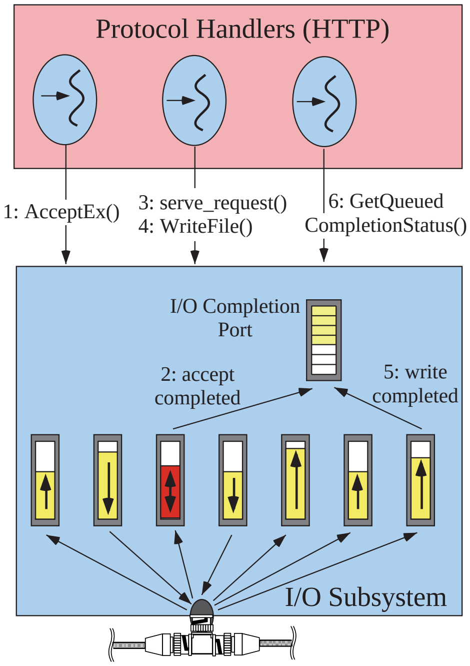

Figure 11: Asynchronous Event Dispatching (Windows NT)

When a JAWS Web server process begins execution, the main thread initiates multiple `AcceptEx` calls asynchronously. All threads in the process block on the completion port by calling `GetQueuedCompletionStatus`. When a new connection arrives, the kernel places the result from `AcceptEx` onto the completion port. If the initial data block received by `AcceptEx` does not contain the entire HTTP request, the selected handler thread must initiate an asynchronous read using `ReadFile`.

After the entire HTTP request is received from the client, the handler thread services the request by locating the file to transmit to the client. If the server is using `WriteFile` to transmit the file data to the client, it memory maps the requested file first.[3] After initiating the asynchronous transfer using `WriteFile` or `TransmitFile`, the handler thread blocks on the completion port. When the asynchronous transfer completes, the result is queued up at the completion port. The OS then selects a thread from those waiting on the completion port. This thread dequeues the result and performs the necessary cleanup (*e.g.*, closes the data socket) to finalize the HTTP transaction.

It is possible that the thread initiating `AcceptEx`, the thread initiating `WriteFile` or `TransmitFile`, and the thread dequeueing the completion status from the port might be different. This increases the adaptivity of JAWS because it is possible to tune the level of concurrency simply by increasing the number of completion handler threads. Moreover, due to the patterns-oriented OO design of the JAWS `Event Dispatcher` framework, these changes do not require any modifications to its concurrency architecture.

4 Benchmarking Testbed and Results

4.1 Hardware Testbed

Our benchmarking hardware testbed is shown in Figure 12. This testbed consists of two Micron Millennia PRO2 plus workstations. Each PRO2 has 128 MB of RAM and is equipped with 2 PentiumPro processors. The client machine has a clock speed of 200 MHz, while the server machine runs 180 MHz. In addition, each PRO2 has an ENI-155P-MF-S ATM card made by Efficient Networks, Inc. and is driven by

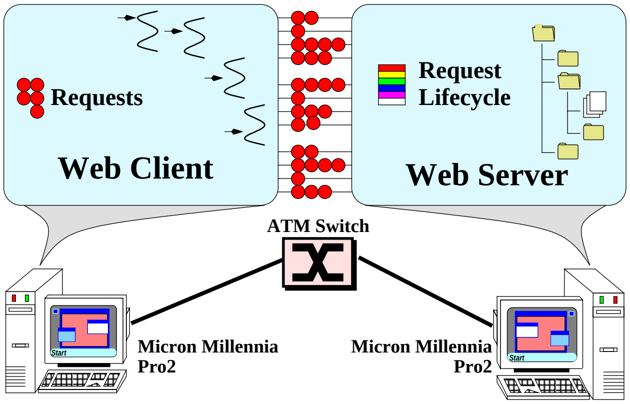

Figure 12: Benchmarking Testbed Overview

---

[3] If the server is using `TransmitFile`, there is no need to memory map the requested file because `TransmitFile` uses the operating system's virtual memory cache manager to retrieve the file data.

Orca 3.01 driver software. The two workstations were connected via an ATM network running through a FORE Systems ASX-200BX, with a maximum bandwidth of 622 Mbps. However, due to limitations of LAN emulation mode, the peak bandwidth of our testbed is approximately 120 Mbps.

4.2 Software Request Generator

We used the WebSTONE [5] v2.0 benchmarking software to collect client- and server-side metrics. These metrics included *average server throughput*, and *average client latency*. WebSTONE is a standard benchmarking utility, capable of generating load requests that simulate typical Web server file access patterns. Our experiments used WebSTONE to generate loads and gather statistics for particular file sizes in order to determine the impacts of different concurrency and event dispatching strategies.

The file access pattern used in the tests is shown in Table 1.

Document Size	**Frequency**
500 bytes	35%
5 Kbytes	50%
50 Kbytes	14%
5 Mbytes	1%

Table 1: File Access Patterns

This table represents actual load conditions on popular servers, based on a study of file access patterns conducted by SPEC [2].

4.3 Experimental Results

The results presented below compare the performance of several different adaptations of the JAWS Web server. We discuss the effect of different event dispatching and I/O models on *throughput* and *latency*, which are measured as follows:

Throughput is defined as the average number of bits received per second by the client. A high-resolution timer for throughput measurement was started before the client benchmarking software sent the HTTP request. The high-resolution timer stops just after the connection is closed at the client end. The number of bits received includes the HTML headers sent by the server.

Latency is defined as the average amount of delay in milliseconds seen by the client from the time it sends the request to the time it completely receives the file. It measures how long an end user must wait after sending an HTTP `GET` request to a Web server, and before the content begins to arrive at the client. The timer for latency measurement is started just before the client benchmarking software sends the HTTP request and stops just after the client receives the first response from the server.

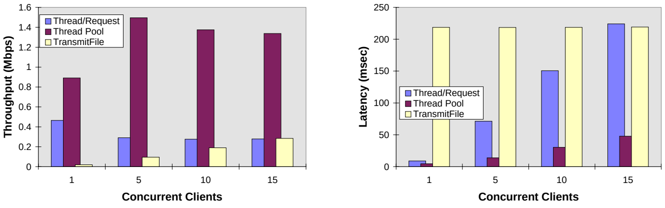

Figure 13: Experiment Results from 500 Byte File

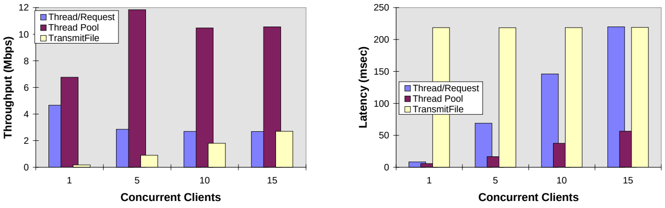

Figure 14: Experiment Results from 5K File

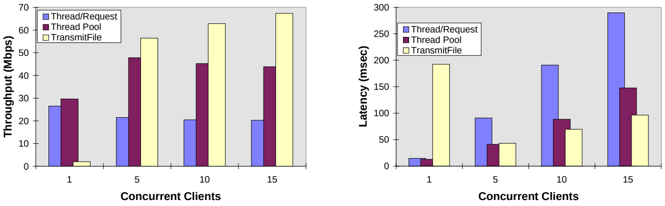

Figure 15: Experiment Results from 50K File

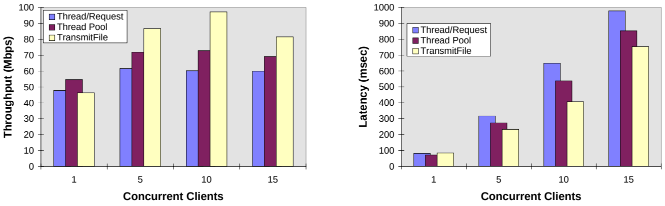

Figure 16: Experiment Results from 500K File

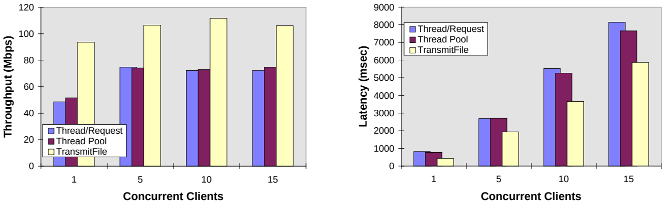

Figure 17: Experiment Results from 5M File

The five graphs shown for each of throughput and latency represent different file sizes used in each experiment, 500 bytes through 5 Mbytes by factors of 10. These files sizes represent the spectrum of files sizes used in our experiments to discover what impact file size has on performance.

4.3.1 Throughput Comparisons

Figures 13-17 demonstrate the variance of throughput as the size of the requested file and the server hit rate are increased systematically. As expected, the throughput for each connection generally degrades as the connections per second increases. This stems from the growing number of simultaneous connections being maintained, which decreases the throughput per connection.

As shown in Figure 15, the throughput of thread-per-request can degrade rapidly for smaller files as the connection load increases. In contrast, the throughput of the synchronous thread pool implementation degrade more gracefully. The reason for this difference is that thread-per-request incurs higher thread creation overhead because a new thread is spawned for each `GET` request. In contrast, thread creation overhead in the thread pool strategy is amortized by pre-spawning threads when the server begins execution.

The results in figures 13-17 illustrate that `TransmitFile` performs poorly for small files (*i.e.,* $<$ 50 Kbytes). Our experiments indicate that the performance of `TransmitFile` depends directly upon the number of simultaneous requests. We believe that during heavy server loads (*i.e.*, high hit rates), `TransmitFile` is forced to wait while the kernel services incoming requests. This creates a high number of simultaneous connections, degrading server performance.

As the size of the file grows, however, `TransmitFile` rapidly outperforms the synchronous dispatching models. For instance, at heavy loads with the 5 Mbyte file (shown in Figure 17), it outperforms the next closest model by nearly 40%. `TransmitFile` is optimized to take advantage of Windows NT kernel features, thereby reducing the number of data copies and context switches.

4.3.2 Latency Comparisons

Figures 13-17 demonstrate the variance of latency performance as the size of the requested file and the server hit rate increase. As expected, as the connections per second increase, the latency generally increases as well. This reflects the additional load placed on the server, which reduces its ability to service new client requests.

As before, `TransmitFile` performs extremely poorly for small files. However, as the file size grows, its latency rapidly improves relative to synchronous dispatching during light loads.

4.3.3 Summary of Benchmark Results

As illustrated in the results presented above, there is significant variance in throughput and latency depending on the concurrency and event dispatching mechanisms. For small files, the synchronous thread pool strategy provides better overall performance. Under moderate loads, the synchronous event dispatching model provides slightly better latency than the asynchronous model. Under heavy loads and with large file transfers, however, the asynchronous model using `TransmitFile` provides better quality of service. Thus, under Windows NT, an optimal Web server should adapt itself to either synchronous or asynchronous event dispatching and file I/O model, depending on the server's workload and distribution of file requests.

5 Comparing JAWS With Other High-performance Web Servers on Windows NT

This section compares a highly optimized JAWS implementation against Netscape Enterprise and Microsoft Internet Information Server (IIS). The results shown in Figures 18-22. demonstrate that JAWS is competitive with state-of-the-art commercial Web server implementations. The figures reveal that JAWS does not perform as well as Enterprise or IIS for small files. However, as the file size grows, JAWS overtakes the other servers in performance.

Two conclusions can be drawn from these results. First, there are still performance issues beyond the scope of this paper that require research to determine how to improve JAWS performance for transferring small files. Second, it affirms our hypothesis that a Web server can only achieve optimal performance by employing adaptive techniques.

6 Related Work

As shown in Figure 1, the JAWS framework focuses on optimizing a range of layers in a Web system. Therefore, it is influenced by a variety of related efforts. This section describes existing work that is most relevant to JAWS.

6.1 Performance Evaluation

The need for high-capacity servers has spurred commercial sector activity and many server implementations are available on the market [27]. The growing number of Web servers has prompted the need for assessing their relative performance. The current standard benchmarks available are WebStone [5] (by SGI) and SPECweb96 [2] (by SPEC), both heavily influenced by the design of LADDIS [31].

WebStone and SPECweb96 are designed to measure overall performance. They rate the performance of a server with a single number (a higher number indicates better performance).

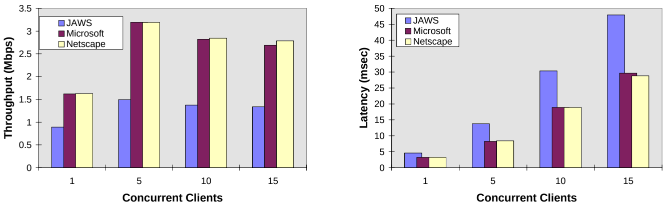

Figure 18: Performance Results from 500 Byte File

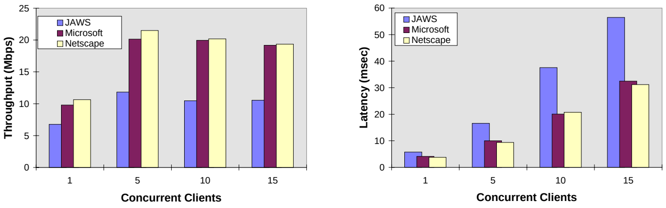

Figure 19: Performance Results from 5K File

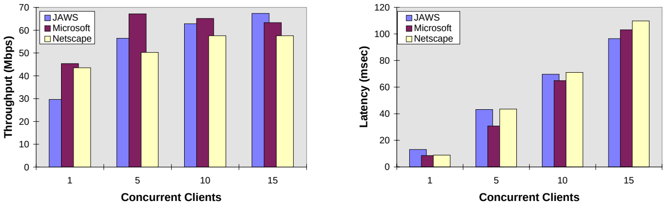

Figure 20: Performance Results from 50K File

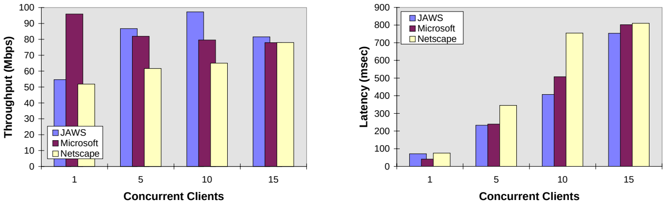

Figure 21: Performance Results from 500K File

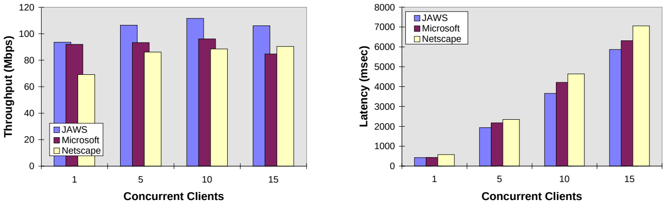

Figure 22: Performance Results from 5M File

These benchmarks are based on a process-based concurrency model and utilize multiple machines to simulate heavy loads. We have found that a thread-based concurrency model allows a single client machine to generate higher loads, thus requiring fewer machines in the benchmarking testbed.

6.2 Optimization Strategies

An important goal of benchmarking Web servers is to discover bottlenecks that require optimization to improve performance. One way to improve performance is by removing overhead in the protocol itself. The W^3C is currently standardizing HTTP/1.1, which multiplexes multiple requests over a single connection. This "connection-caching" strategy can significantly enhance the performance over HTTP/1.0 by reducing unnecessary connection set up and termination [28, 20]. The need for persistent connections to improve latency has been noted in [15].

Latency can be improved by using caching proxies and caching clients, although the removal policy must be considered carefully [30]. The efforts presented in this paper are orthogonal to work on client-side caching. Our attempts to optimize performance have focused on (1) server-side caching, *e.g.*, via JAWS' Cached Virtual File system and (2) utilizing concurrency and event dispatching processing routines that are customized for the OS platform to reduce server load and improve end-to-end quality of service.

6.3 Web Server Design and Implementation

Design patterns are re-usable software abstractions that have been observed to occur many times in actual solutions [4]. A design pattern is intended to solve a general design problem for a specific context. Many patterns have been observed in the context of concurrent, parallel and distributed systems. Many of these ideas are applicable to Web server design.

Katz presents an NCSA prototype of a scalable Web server design [11]. This design was prompted by the growing number of requests being made to the NCSA server machine. Many commercial server implementations arose to meet the demand for high-performance Web servers. Higher end implementations, such as Netscape Enterprise and Zeus, use multi-threading to scale for high-end architectures with multiple CPUs. Other implementations (*e.g.*, Roxen, BOA and thttpd) use a single thread of control to optimize performance on single CPU architectures. Yeager and McGrath of NCSA discuss the tradeoffs of some Web server designs and their performance in [17].

Our work attempts to apply design patterns and concurrency strategies to provide a design which is both scalable and adaptive. We achieve static adaptation through start-up time configuration options, and dynamic adaptation through use of dynamically loadable modules which alter the behavior of the server.

6.4 Concurrency Strategies

A concurrency strategy specifies the policies and mechanisms for allowing multiple tasks to execute simultaneously. For Web servers, a task is an object that encapsulates server request handling. There are four general classes of Web server concurrency strategies:

Iterative: An iterative server is a server that has no application-level concurrency: each task that begins runs to completion.

Single-threaded concurrent: A single-threaded concurrent server has one thread of control, but only partially serves each unfinished task until a pre-determined amount of work is done, or until a timer has expired.

Thread-per-Request: The thread-per-request strategy allows each task to execute in its own thread of control.

Thread Pool: A thread pool approach fixes the number of executing threads and assigns tasks to available threads.

Surprisingly, Web server literature contains relatively little information on the performance of alternative concurrency strategies. This may result from the fact that most Web server implementations tightly couple their concurrency strategy with the other components in the server. In contrast, concurrency strategies have been studied extensively in the context of parallel protocol stacks that run on shared memory platforms.

[13] measured the impact of synchronization on thread-per-request implementations of TCP and UDP transport protocols built within a multi-processor version of the *x*-kernel; [16] examined performance issues in parallelizing TCP-based and UDP-based protocol stacks using a thread-per-request strategy in a different multi-processor version of the *x*-kernel; and [21] measured the performance of the TCP/IP protocol stack using a thread-per-connection strategy in a multi-processor version of System V STREAMS.

The results presented in this paper extend existing research on protocol stack parallelism in several ways. First, we measure the performance of a variety of representative concurrency strategies. Second, our experiments report the impact of synchronous and asynchronous event dispatching strategies on Web server performance, whereas existing work on parallel protocol stacks has focused on synchronous event dispatching. Our results illustrate how operating systems like Windows NT that support asynchronous event dispatching strategies and customized file/network transfer operations can perform significantly better than purely synchronous strategies under various workloads (*e.g.*, transfers involving large files).

7 Concluding Remarks

This paper describes *static* and *dynamic* adaptations that can be applied to develop high-performance Web servers. Common static adaptations include configuring an event dispatching model that is customized for OS platform-specific features, such as the I/O completion ports and `TransmitFile` on Windows NT. Common dynamic adaptations include prioritized request handling, caching strategies, and threading strategies. We illustrate the results of adapting the JAWS Web server framework to leverage the native asynchronous dispatching mechanisms provided by Windows NT.

Our results demonstrate that to alleviate bottlenecks, Web servers must utilize an integrated approach that combines optimizations at multiple levels of a Web endsystem. For example, a Web server should take advantage of special I/O system calls, specialized hardware and knowledge of the file access patterns.

As illustrated in Section 4, there is significant variance in throughput and latency under different server load conditions. For small files, the synchronous thread pool strategy provides better overall performance. However, under heavy loads and with large file transfers, the asynchronous model using `TransmitFile` provides more consistent quality of service. Thus, under Windows NT, an optimal Web server should adapt itself dynamically to either event dispatching and file I/O model, depending on the server's workload.

Acknowledgements

The authors thank the reviewers for their thoughtful comments and suggestions for improvement. We would also like to thank Chris Dawson and Dan Swanger of FORE Systems, Inc., for providing the ATM cards and ATM switch used for our benchmarking testbed. In addition, we would like to thank Sumedh Mungee, Darrell Brunsch, Everett Anderson, Tim Harrison, and John DeHart for helping to develop our ATM Web server testbed.

References

[1] Frank Buschmann, Regine Meunier, Hans Rohnert, Peter Sommerlad, and Michael Stal. *Pattern-Oriented Software Architecture - A System of Patterns*. Wiley and Sons, 1996.

[2] Alexander Carlton. An Explanation of the SPECweb96 Benchmark. Standard Performance Evaluation Corporation whitepaper, 1996. www.specbench.org/.

[3] Zubin D. Dittia, Guru M. Parulkar, and Jerome R. Cox, Jr. The APIC Approach to High Performance Network Interface Design: Protected DMA and Other Techniques. In *Proceedings of INFOCOM '97*, Kobe, Japan, April 1997. IEEE.

[4] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, Reading, MA, 1995.

[5] Gene Trent and Mark Sake. WebSTONE: The First Generation in HTTP Server Benchmarking. Silicon Graphics, Inc. whitepaper, February 1995. ww.sgi.com/.

[6] James Hu, Sumedh Mungee, and Douglas C. Schmidt. Principles for Developing and Measuring High-performance Web Servers over ATM. In *Proceeedings of INFOCOM '98*, March/April 1998.

[7] James Hu, Irfan Pyarali, and Douglas C. Schmidt. Measuring the Impact of Event Dispatching and Concurrency Models on Web Server Performance Over High-speed Networks. In *Proceedings of the* 2^{nd} *Global Internet Conference*. IEEE, November 1997.

[8] James Hu, Irfan Pyarali, and Douglas C. Schmidt. Applying the Proactor Pattern to High-Performance Web Servers. In *Proceedings of the 10th International Conference on Parallel and Distributed Computing and Systems*. IASTED, October 1998.

[9] Prashant Jain and Douglas C. Schmidt. Dynamically Configuring Communication Services with the Service Configurator Pattern. *C++ Report*, 9(5), June 1997.

[10] Prashant Jain and Douglas C. Schmidt. Service Configurator: A Pattern for Dynamic Configuration of Services. In *Proceedings of the* 3^{rd} *Conference on Object-Oriented Technologies and Systems*. USENIX, June 1997.

[11] Eric Dean Katz, Michelle Butler, and Robert McGrath. A Scalable HTTP Server: The NCSA Prototype. In *Proceedings of the First International Conference on the World-Wide Web*, May 1994.

[12] R. Greg Lavender and Douglas C. Schmidt. Active Object: an Object Behavioral Pattern for Concurrent Programming. In James O. Coplien, John Vlissides, and Norm Kerth, editors, *Pattern Languages of Program Design*. Addison-Wesley, Reading, MA, 1996.

[13] Mats Bjorkman and Per Gunningberg. Locking Strategies in Multiprocessor Implementations of Protocols. *Transactions on Networking*, 3(6), 1996.

[14] *Microsoft Developers Studio, Version 4.2 - Software Development Kit*, 1996.

[15] Jeffrey C. Mogul. The Case for Persistent-connection HTTP. In *Proceedings of ACM SIGCOMM '95 Conference in Computer Communication Review*, pages 299–314, Boston, MA, USA, August 1995. ACM Press.

[16] Erich M. Nahum, David J. Yates, James F. Kurose, and Don Towsley. Performance Issues in Parallelized Network Protocols. In *Proceedings of the* 1^{st} *Symposium on Operating Systems Design and Implementation*. USENIX Association, November 1994.

[17] Nancy J. Yeager and Robert E. McGrath. *Web Server Technology: The Advanced Guide for World Wide Web Information Providers*. Morgan Kaufmann, 1996.

[18] Object Management Group. *Control and Management of Audio/Video Streams: OMG RFP Submission*, 1.2 edition, March 1997.

[19] Irfan Pyarali, Timothy H. Harrison, and Douglas C. Schmidt. Asynchronous Completion Token: an Object Behavioral Pattern for Efficient Asynchronous Event Handling. In Robert Martin, Frank Buschmann, and Dirk Riehle, editors, *Pattern Languages of Program Design*. Addison-Wesley, Reading, MA, 1997.

[20] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, and T. Berners-Lee. Hypertext Transfer Protocol – HTTP/1.1. Standards Track RFC 2068, Network Working Group, January 1997. www.w3.org/.

[21] Sunil Saxena, J. Kent Peacock, Fred Yang, Vijaya Verma, and Mohan Krishnan. Pitfalls in Multithreading SVR4 STREAMS and other Weightless Processes. In *Proceedings of the Winter USENIX Conference*, pages 85–106, San Diego, CA, January 1993.

[22] Douglas C. Schmidt. GPERF: A Perfect Hash Function Generator. In *Proceedings of the* 2^{nd} *C++ Conference*, pages 87–102, San Francisco, California, April 1990. USENIX.

[23] Douglas C. Schmidt. ACE: an Object-Oriented Framework for Developing Distributed Applications. In *Proceedings of the* 6^{th} *USENIX C++ Technical Conference*, Cambridge, Massachusetts, April 1994. USENIX Association.

[24] Douglas C. Schmidt. Reactor: An Object Behavioral Pattern for Concurrent Event Demultiplexing and Event Handler Dispatching. In James O. Coplien and Douglas C. Schmidt, editors, *Pattern Languages of Program Design*, pages 529–545. Addison-Wesley, Reading, MA, 1995.

[25] Douglas C. Schmidt. Acceptor and Connector: Design Patterns for Initializing Communication Services. In Robert Martin, Frank Buschmann, and Dirk Riehle, editors, *Pattern Languages of Program Design*. Addison-Wesley, Reading, MA, 1997.

[26] Douglas C. Schmidt and James Hu. Developing Flexible and High-performance Web Servers with Frameworks and Patterns. *ACM Computing Surveys*, 30, 1998.

[27] David Strom. Web Compare. webcompare.iworld.com/, 1997.

[28] T. Berners-Lee, R. T. Fielding, and H. Frystyk. Hypertext Transfer Protocol – HTTP/1.0. Informational RFC 1945, Network Working Group, May 1996. www.w3.org/.

[29] Stephen Williams, Marc Abrams, Charles R. Standridge, Ghalleb Abdulla, and Edward A. Fox. Removal Policies in Network Caches for World Wide Web Documents. In *Proceedings of SIGCOMM '96*, pages 293–305, Stanford, CA, August 1996. ACM.

[30] Stephen Williams, Marc Abrams, Charles R. Standridge, Ghalleb Abdulla, and Edward A. Fox. Removal Policies in Network Caches for World Wide Web Documents. In *Proceedings of SIGCOMM '96*, pages 293–305, Stanford, CA, August 1996. ACM.

[31] Mark Wittle and Bruce E. Keith. LADDIS: The Next Generation in NFS File Server Benchmarking. In *USENIX Association Conference Proceedings '93*, April 1993.