

Benefits of I/O Acceleration Technology (I/OAT) in Clusters*

Karthikeyan Vaidyanathan

Dhabaleswar K. Panda

Computer Science and Engineering,

The Ohio State University

{vaidyana, panda}@cse.ohio-state.edu

Abstract

Packet processing in the TCP/IP stack at multi-Gigabit data rates occupies a significant portion of the system overhead. Though there are several techniques to reduce the packet processing overhead on the sender-side, the receiver-side continues to remain as a bottleneck. I/O Acceleration Technology (I/OAT), developed by Intel, is a set of features particularly designed to reduce the receiver-side packet processing overhead. This paper studies the benefits of the I/OAT technology by extensive evaluations through micro-benchmarks as well as evaluations on two different application domains: (1) A multi-tier data-center environment and (2) A Parallel Virtual File System (PVFS). Our micro-benchmark evaluations show that I/OAT results in 38% lower overall CPU utilization in comparison with traditional communication. Due to this reduced CPU utilization, I/OAT delivers better performance and increased network bandwidth. Our experimental results with data-centers and file systems reveal that I/OAT can improve the total number of transactions processed by 14% and throughput by 12%, respectively. In addition, I/OAT can sustain a large number of concurrent threads (up to a factor of four as compared to non-I/OAT) in data-center environments, thus increasing the scalability of the servers.

1 Introduction

Over the past few years, there has been an incredible growth of highly data-intensive applications in various fields such as medical informatics, genomics, e-commerce, data mining and satellite weather image analysis. With technology trends, the ability to store and share the datasets generated by these applications is also increasing, allowing scientists and institutions to create large dataset repositories and making them available for use by others. On the other hand, clusters consisting of commodity off-the-shelf hardware components have become increasingly attractive as platforms for high-performance computation and scalable servers. Based on these two trends, researchers have proposed the feasibility and potential of cluster-based servers [14, 10, 18, 19].

Several clients request these servers for either the raw or some kind of processed data simultaneously. However, existing servers are becoming increasingly incapable of

meeting such sky-rocketing processing demands with high-performance and scalability. These servers rely on TCP/IP for data communication and typically use Gigabit Ethernet networks for cost-effective designs. The host-based TCP/IP protocols on such networks have high CPU utilization and low bandwidth, thereby limiting the maximum capacity (in terms of requests they can handle per unit time). Alternatively, many servers use multiple Gigabit Ethernet networks to cope with the network traffic. However, at multi-Gigabit data rates, packet processing in the TCP/IP stack occupies a significant portion of the system overhead.

Packet processing [12, 13] usually involves manipulating the headers and moving the data through the TCP/IP stack. Though this does not require significant computation, processor time gets wasted due to delays caused by latency of memory accesses and data movement operations. To overcome these overheads, researchers have proposed several techniques [9] such as transport segmentation offload (TSO), jumbo frames, zero-copy data transfer (sendfile()), interrupt coalescing, etc. Unfortunately, many of these techniques are applicable only on the sender side, while the receiver side continues to remain as a bottleneck in several cases, thus resulting in a huge performance gap between the CPU overheads of sending and receiving packets.

Intel's I/O Acceleration Technology (I/OAT) [1, 3, 2, 15] is a set of features which attempts to alleviate the receiver packet processing overheads. It has three additional features, namely: (i) split headers, (ii) DMA copy offload engine and (iii) multiple receive queues.

At this point, the following open questions arise:

- What kind of benefits can be expected from the current I/OAT architecture?
- How does this benefit translate to applications?

In this paper, we focus on the above questions. We first analyze the performance of I/OAT based on a detailed suite of micro-benchmarks. Next, we evaluate it on two different application domains:

- A multi-tier Data-Center environment
- A Parallel Virtual File System (PVFS)

Our micro-benchmark evaluations show that I/OAT reduces the overall CPU utilization significantly, up to 38%, as compared

*This research is supported in part by NSF grants #CNS-0403342 and #CNS-0509452; DOE grants #DE-FC02-06ER25749 and #DE-FC02-06ER25755; grants from Intel, Mellanox, Cisco systems, Linux Network and Sun Microsystems; and equipment donations from Intel, Mellanox, AMD, Apple, Appro, Dell, Microway, PathScale, IBM, SilverStorm and Sun Microsystems.

to traditional communication (non-I/OAT). Due to this reduced CPU utilization, I/OAT delivers better performance and increased network bandwidth. Our experimental results with data-centers and file systems reveal that I/OAT can improve the total number of transactions processed by 14% and the throughput by 12%, respectively. Also, our results show that I/OAT can sustain a large number of concurrent threads (up to a factor of four as compared to non-I/OAT) in data-center environments, thus increasing the scalability of the servers.

The remaining part of the paper is organized as follows: Section 2 provides a brief background on socket optimizations and I/OAT architecture. In Section 3, we discuss the software infrastructures used; in particular the Multi-Tier Data-Center environments and Parallel Virtual File System (PVFS). Section 4 addresses micro-benchmark evaluations to study the ideal case benefits of I/OAT over the native kernel implementation. Sections 5 and 6 present the benefits of I/OAT in a data-center and PVFS environment, respectively. Section 7 presents the advantages and the disadvantages of I/OAT. We conclude the paper in Section 8.

2 Background

In this section, we present some of the socket optimization technologies in detail. Later, we provide a brief background of the I/OAT architecture and its features.

2.1 Socket Optimizations

As mentioned in Section 1, due to technological developments, significant performance gaps exist between the CPU overhead of sending and receiving packets. In particular, there are two techniques that make the CPU usage on the sender-side far less than the CPU usage on the receiver-side.

The first technique, TCP segmentation offload (TSO), allows the kernel to send a large buffer, which is more than the maximum transmission unit (MTU), to the network controller. The network controller segments the buffer into individual Ethernet frames and sends it on the wire. In the absence of this technique, the host CPU is expected to break the large buffer into small frames and send the small frames to the network controller. This operation is more CPU intensive. The second technique, *sendfile()* optimization, allows the kernel to perform zero-copy operation. Using this technique, the kernel does not copy the user buffer to the network buffer. Instead, it points the pinned pages of the user buffer as the data source for transmission.

However, neither of these techniques are applicable to the receiver side. On the receiver side, interrupt coalescing technique helps to reduce the number of interrupts generated in the host system. This technique generates one interrupt for multiple packets rather than one interrupt for every single packet. However, this optimization improves the performance only when the network is heavily loaded.

2.2 I/O Acceleration Technology (I/OAT) Overview

I/OAT [15] attempts to alleviate the bottlenecks mentioned in Section 1 by providing a set of features, namely: (i) Split

headers, (ii) Asynchronous copy using DMA engine and (iii) Multiple Receive Queues.

2.2.1 Split-headers

As shown in Figure 1, the optimized TCP/IP stack of I/OAT has the split header feature implemented. In TCP/IP-based communication for transmission of application data, headers such as TCP header, IP header and Ethernet header are attached along with the application data. Typically, the network controller transfers the headers and the application data onto a single buffer. However, with the split-header feature, the controller partitions the network data into headers and application data and copies them into two separate buffers. This allows the header and data to be optimally aligned. Since the headers are frequently accessed during network processing, this feature also results in better cache utilization by not polluting the CPU’s cache with any application data while accessing the headers, thus increasing the locality of the incoming headers. For more information regarding this feature and its benefits please refer to [15, 16].

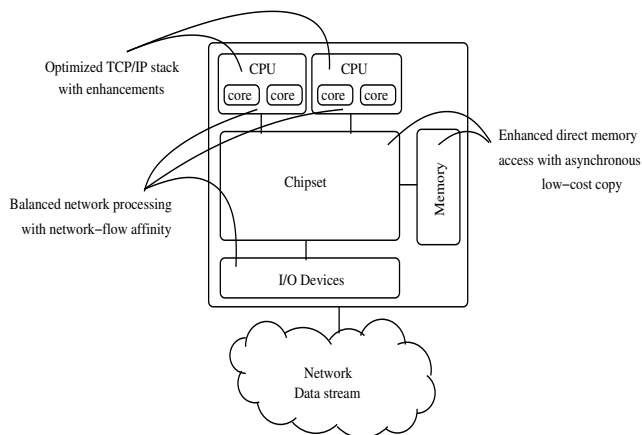


Figure 1. Intel’s I/O Acceleration Technology (Courtesy Intel)

2.2.2 Asynchronous copy using DMA engine

Most of the receive packet processing time [15] is spent in copying the data from kernel buffer to user buffer. In particular, the ratio of useful work done by the CPU to useless overheads such as CPU stalling for memory access or data movement, decreases as we go from 1 Gbps links to 10 Gbps links [17]. Especially for data movement operations, rather than waiting for a memory copy to finish, the host CPU can process other pending packets while the copy is still in progress. I/OAT offloads this copy operation with an additional DMA engine. This is a dedicated device which can perform memory copies. As a result, while the DMA is performing the data copy, the CPU becomes free to process other pending packets. For more details regarding this feature and its implementation please refer to [15, 17, 16].




Figure 2. Application Domains: (a) A Multi-Tier Data-Center Environment and (b) Parallel Virtual File System (PVFS) Environment

2.2.3 Multiple Receive Queues

Processing large packets is generally not CPU-intensive, whereas processing small packets can fully occupy the CPU. Even on a multi-CPU system, processing occurs on a single CPU (the CPU which handles the controller's interrupt). Multiple receive queue feature helps in distributing the packets among different receive queues and allows multiple CPUs to act on different packets at the same time. For more details regarding this feature please refer to [15, 16]. Unfortunately, this feature is currently disabled in the Linux platform. Hence, we could not measure the performance impact of this feature in our evaluation.

3 Software Infrastructure

We have carried out the evaluation of I/OAT on two different software infrastructures: Multi-Tier Data Center environment and the Parallel Virtual File System (PVFS). In this section, we discuss each of these in more detail.

3.1 Multi-Tier Data Center environment

More and more people are using web interfaces for a wide range of services. Scalability of these systems is a very important requirement. Upgrading a server to a single, more powerful server is no longer a feasible method. In recent years, clusters consisting of low-cost commodity off-the-shelf hardware components have become increasingly attractive for hosting web services. Cluster architectures are characterized by a low setup and maintenance cost. In addition, clusters provide greater flexibility in reallocating resources among the tiers of a data-center to accommodate fluctuating load.

A typical Multi-tier Data-center [20, 8, 21] has a cluster of nodes, called the edge nodes, as its first tier. These nodes can be thought of as switches (up to the 7th layer) that provide services like load balancing, security, caching etc. The main purpose of this tier is to increase the performance of the inner tiers. The next tiers are usually the web-servers. These nodes, apart from serving static content, can fetch dynamic data from other sources and serve that data in a presentable form. The last tier of the data-center is the database tier. It is used to store persistent data. This tier is usually I/O intensive. Figure 2a shows a typical data-center setup.

A request from a client is received by the edge or proxy servers. This request is serviced from cache if possible, otherwise it is forwarded to the web/application servers. Static requests are serviced by the web servers by just returning the requested file to the client via the edge server. This content may be cached at the edge server so that subsequent requests to the same static content may be served from the cache. The application tier handles dynamic content. Any request that needs some value to be computed, searched, analyzed or stored, must use this tier at some stage. The application servers may also need to spend some time on converting data to presentable formats. The back-end database servers are responsible for storing data persistently and responding to queries. These nodes are connected to a persistent storage system. Queries to the database systems can be anything from a simple search of required data to performing join/aggregation/select operations on the data.

3.2 Parallel Virtual File System (PVFS)

Parallel Virtual File System (PVFS) [7] is one of the leading parallel file systems for Linux clusters today. It is designed to meet the increasing I/O demands of parallel applications in cluster systems. Figure 2b demonstrates a typical PVFS environment. As shown in the figure, a group of nodes in the cluster system can be configured as I/O servers and one of them (either an I/O server or a different node) as a meta-data manager. It is possible for a node to host computation while serving as an I/O node.

PVFS achieves high performance by striping files across a set of I/O server nodes to allow parallel accesses to the data. It uses the native file system on the I/O servers to store individual file stripes. An I/O daemon runs on each I/O node and services requests from the compute nodes, in particular the read and write requests. Data is transferred directly between the I/O servers and the compute nodes.

A manager daemon runs on the meta-data manager node. It handles meta-data operations involving file permissions, truncation, file stripe characteristics, and so on. Meta-data is also stored on the local file system. The meta-data manager provides a cluster-wide consistent name space to applications. In PVFS, the meta-data manager does not participate in read/write operations.

PVFS supports a set of feature-rich interfaces, including


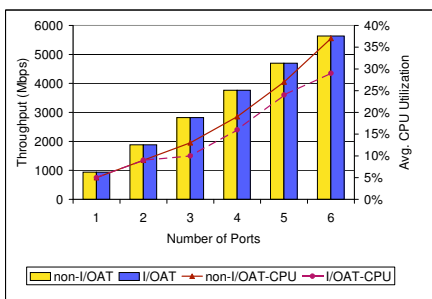


Figure 3. Micro-Benchmarks: (a) Bandwidth and (b) Bi-directional Bandwidth

support for both contiguous and non-contiguous accesses to memory and files [11]. PVFS can be used with multiple APIs: a native API, the UNIX/POSIX API, MPI-IO [22], and an array I/O interface called the Multi-Dimensional Block Interface (MDBI). The presence of multiple popular interfaces contributes to the wide success of PVFS in the industry.

4 I/OAT Micro-Benchmark Results

In this section, we compare the ideal case performance benefits achievable by I/OAT as compared to the native sockets implementation (non-I/OAT) using a set of micro-benchmarks.

We use two testbeds for all of our experiments. Their descriptions are as follows:

Testbed 1: A system consisting of two nodes built around SuperMicro X7DB8+ motherboards which include 64-bit 133 MHz PCI-X interfaces. Each node has a dual dual-core Intel 3.46 GHz processor with a 2 MB L2 cache. The machines are connected with three Intel PRO 1000Mbit adapters with two ports each through a 24-port Netgear Gigabit Ethernet switch. We use the Linux RedHat AS 4 operating system and kernel version 2.6.9-30.

Testbed 2: A cluster system consisting of 44 nodes. Each node has a dual Intel Xeon 2.66 GHz processor with 512KB L2 cache and 2GB of main memory.

For the experiments mentioned in Sections 5, we use the nodes in Testbed 2 as clients and the nodes in Testbed 1 as servers. For all other experiments, we only use the nodes in Testbed 1. Also, for experiments within Testbed 1, we create a separate VLAN for each network adapter in one node and a corresponding IP address within the same VLAN on the other node to ensure an even distribution of network traffic. In all of our experiments, we define the term relative CPU benefit of I/OAT as follows: if a is the % CPU utilization of I/OAT and b is the % CPU utilization of non-I/OAT, the relative CPU benefit of I/OAT is defined as $(b - a)/b$. For example, if I/OAT occupies 30% CPU and non-I/OAT occupies 60% CPU, the relative CPU benefit of I/OAT is 50%, though the absolute difference in CPU usage is only 30%.

4.1 Bandwidth and Bi-directional Bandwidth

Figure 3a shows the bandwidth achieved by I/OAT and non-I/OAT with an increasing number of network ports. We use the standard *tcp* benchmark [4] for measuring the bandwidth. As the number of ports increase, we expect the bandwidth to increase. As shown in Figure 3a, we see that the bandwidth

performance achieved by I/OAT is similar to the performance achieved by non-I/OAT with an increasing number of ports. The maximum bandwidth achieved is close to 5635 Mbps with six network ports. However, we see a difference in performance with respect to the CPU utilization on the receiver side. We observe that the CPU utilization is lower for I/OAT as compared to non-I/OAT using three network ports and this difference increases as we see the number of ports increase from three to six. For a six port configuration, non-I/OAT occupies close to 37% of the CPU while I/OAT occupies only 29% of the CPU. The relative benefit achieved by I/OAT in this case is close to 21%.

In the bi-directional bandwidth test, we use two machines and $2*N$ threads on each machine with N threads acting as servers and the other N threads as clients. Each thread on one machine has a connection to exactly one thread on the other machine. The client threads connect to the server threads on the other machine. Thus, $2*N$ connections are established between these two machines. On each connection, the basic bandwidth test is performed using the *ttcp* benchmark. The aggregate bandwidth achieved by all threads is calculated as the bi-directional bandwidth, as shown in Figure 3b. In our experiments, N is equal to the number of network ports. We see that the maximum bi-directional bandwidth is close to 9600 Mbps. Also, we observe that I/OAT shows an improvement in CPU utilization using only two ports and this improvement increases with an increasing number of ports. With six network ports, non-I/OAT occupies close to 90% of CPU whereas I/OAT occupies only 70% of the CPU. The relative CPU benefits achieved by I/OAT is close to 22%. This trend also suggests that with an addition of one or two network ports to this configuration, non-I/OAT may not give the best network throughput in comparison with I/OAT since non-I/OAT may end up occupying 100% CPU.

4.2 Multi-Stream Bandwidth

The multi-stream bandwidth test is very similar to the bi-directional bandwidth test mentioned above. However, in this experiment, only one machine acts as a server and the other machine as the client. We use two machines and N threads on each machine. Each thread on one machine has a connection to exactly one thread on the other machine. On each connection, the basic bandwidth test is performed. The aggregate bandwidth achieved by all threads is calculated as the multi-stream bandwidth. As shown in Figure 4, we observe that the band-

width achieved by I/OAT is similar to the bandwidth achieved by non-I/OAT for an increasing number of threads. However, when the number of threads increases to 120, we see a degradation in performance of non-I/OAT; whereas, I/OAT consistently shows no degradation in network bandwidth. Further, the CPU utilization for non-I/OAT also increases with an increasing number of threads. With 120 threads in the system, we see that non-I/OAT occupies close to 76% CPU whereas I/OAT only occupies 52% resulting in 24% absolute benefit in CPU utilization. The relative CPU benefits achieved by I/OAT in this case is close to 32%.

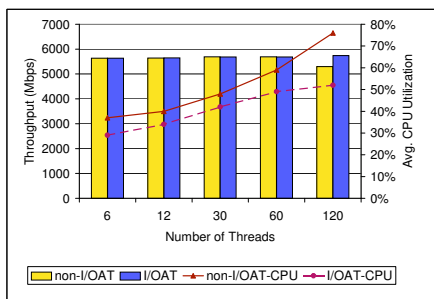


Figure 4. Multi-Stream Bandwidth

4.3 Bandwidth and Bi-directional Bandwidth with Socket Optimizations

As mentioned in Section 2, several optimizations on the sender side exist to reduce the packet overhead and also to improve the network performance. In this experiment, we considered three such existing optimizations: (i) Large Socket buffer sizes (100 MB), (ii) Segmentation Offload (TSO) and (iii) Jumbo Frames. In this experiment, we aim to understand the impact of each of these optimizations and observe the improvement, both in terms of throughput and CPU utilization on the receiver-side. Case 1 uses the default socket options without any optimization. In Case 2, we increase the socket buffer size to 100 MB. For Case 3, we further improve the optimization by enabling segmentation offload (TSO) so that the host CPU is relieved from fragmenting large packets. In Case 4, in addition to the previous optimizations, we increase the MTU-size to 2048 bytes so that large packets are sent over the network. In addition to the sender-side optimizations, in Case 5, we include the interrupt coalescing feature. Performance numbers with various socket optimizations are shown in Figure 5.

In the bandwidth test, as shown in Figure 5a, we observe two interesting trends. First, as we increase the socket optimizations, we see an increase in the aggregate bandwidth. Second, we observe that the performance of I/OAT is consistently better than the performance of non-I/OAT. Especially for Case 5, which includes all socket optimizations, the bandwidth achieved by I/OAT is close to 5586 Mbps, whereas non-I/OAT achieves only 5514 Mbps. More importantly, we observe that there is a significant improvement in terms of CPU utilization. As shown in Figure 5a, the relative CPU benefit of I/OAT increases as we increase the socket optimizations. Especially

in Case 4, we observe that I/OAT achieves 30% relative CPU benefit in comparison with non-I/OAT.

We see similar trends with the bi-directional bandwidth test as shown in Figure 5b. Further, the relative benefits achieved by I/OAT is much higher than compared to the bandwidth experiment. In Case 4, I/OAT achieves close to 38% relative CPU benefit as compared to non-I/OAT.

In summary, we see that the micro-benchmark results show a significant improvement in terms of CPU utilization and network performance for I/OAT. Since I/OAT in Linux has two additional features, the split-header and DMA copy engine, it is also important to understand the benefits attributed by each of these features individually. In the following section, we conduct several experiments to show the individual benefits.

4.4 Benefits of Asynchronous DMA Copy Engine

In this section, we isolate the DMA engine feature of I/OAT and show the benefits of an asynchronous DMA copy engine. We compare the performance of the copy engine with that of traditional CPU-based copy and show its benefits in terms of performance and overlap efficiency. For CPU-based copy, we use the standard *memcpy* utility.

Figure 6 compares the cost of performing a copy using the CPU and I/OAT’s DMA engine. The *copy-cache* bars denote the performance of CPU-based copy with the source and destination buffers in the cache and the *copy-nocache* bars denote the performance with the source and destination buffers not in the cache. The *DMA-copy* bars denote the total time taken to perform the copy using the copy engine and the *DMA-overhead* bars include the startup overhead in initiating the copy using the DMA engine. The *Overlap* line denotes the percentage of DMA copy time that can be overlapped with other computations. As shown in Figure 6, we see that the performance of DMA-based copy approach (*DMA-copy*) is better than the performance of CPU-based copy approach (*copy-nocache*) for message sizes greater than 8 KB. Further, we observe that the percentage of overlap increases with increasing message sizes reaching up to 93% for a 64 KB message size. However, if the source and destination buffers are in the cache, we observe that the performance of the CPU-based copy is much better than the performance of the DMA-based copy approach. Also, since the DMA-based copy can be overlapped with processing other packets, we incur only the DMA startup overheads. As observed in Figure 6, we see that the DMA startup overhead time is much less than the time taken by the CPU-based copy approach. Thus, the DMA copy engine can also be useful even if the source and destination buffers are in cache.

4.5 Breakdown of I/OAT Features

In this section, we break down the two features of I/OAT and show its benefits in terms of throughput and CPU utilization. We first measure the performance of non-I/OAT without the DMA engine and the split header features. Next, we measure the performance including the DMA engine feature (denoted as I/OAT-DMA in Figure 7) and compare it with the

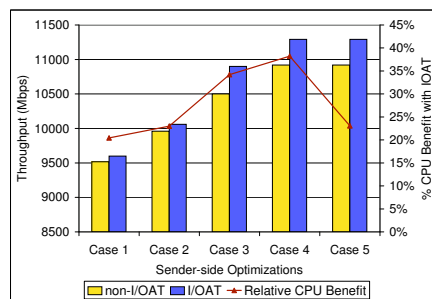



Figure 5. Optimizations: (a) Bandwidth and (b) Bi-directional Bandwidth

performance of non-I/OAT. We report the difference in performance as DMA engine benefits. Finally, we measure the performance including both the split-header feature and the DMA engine feature (denoted as I/OAT-SPLIT in Figure 7) and compare it with the performance of I/OAT with DMA engine. We report the difference in performance as split-header benefits.

We design the experiment in the following way. We use the two nodes from Testbed 1, one node as a server and the other as a client. We use two Intel network adapters with two ports each on both the server and the client node. Since we have four network ports on each node, we use four clients and perform a bandwidth test with four server threads. We first measure the bandwidth performance for small message sizes for all three cases, i.e., non-I/OAT, I/OAT-DMA and I/OAT-SPLIT and report the relative CPU benefits. As shown in Figure 7a, we observe that the DMA engine feature achieves close to 16% relative CPU benefit compared to the non-I/OAT case. This improvement is seen for all message sizes from 16 KB to 128 KB. However, we do not observe any throughput improvement for all message sizes. Also, the split-header feature of I/OAT does not seem to improve the CPU or the throughput for all message sizes.

As mentioned in Section 2, the split-header feature of I/OAT increases the locality of incoming headers and also helps improve cache utilization by not polluting the CPU’s cache with application buffers during network processing. In order to show the cache pollution effects, we repeat the previous experiment for large message sizes that cannot fit in the system cache. Figure 7b shows the throughput benefits achieved by I/OAT features for large message sizes. We observe that the split-header feature can achieve up to 26% benefit in throughput for transferring 1 MB of application data. In this case, the server with I/OAT capability receives a total of 4 MB of application data from the four clients and since the cache size is only 2 MB, clearly the application data does not fit in the system cache. We see a huge improvement for the split-header feature because the CPU’s cache is not polluted with application data. In addition, we see that the benefits achieved by the split-header feature decreases for increasing message sizes.

5 Data-Center Performance Evaluation

In this section, we analyze the performance of a 2-tier data-center environment with I/OAT and compare its performance with non-I/OAT. For all experiments in this section, we use the nodes in Testbed 1 (described in Section 4) for the data-center




Figure 6. Copy Performance: CPU vs DMA

tiers. For the client nodes, we use the nodes in Testbed 2 for most of the experiments. We notify the readers at appropriate points in the remaining sections when other nodes are used as clients.

5.1 Evaluation Methodology

As mentioned in Section 2, I/OAT is a server architecture geared to improve the receiver-side performance. In other words, I/OAT can be deployed in a data-center environment and client requests coming over the Wide Area Network (WAN) can seamlessly take advantage of I/OAT and get serviced much faster. Further, the communication between the tiers inside the data-center, such as proxy tier and application tier, can be greatly enhanced using I/OAT, thus improving the overall data-center performance and scalability.

We set up a two-tier data-center testbed to determine the performance characteristics of using I/OAT and non-I/OAT. The first tier consists of the front-end proxies. For this, we used the proxy module of Apache 2.2.0. The second tier consists of the web server module of Apache in order to service static requests. The two tiers in the data-center reside on 1 Gigabit Ethernet network; the clients are connected to the data-center using a 1 Gigabit Ethernet connection.

Typical data-center workloads have a wide range of characteristics. Some workloads may vary from high to low temporal locality, following a Zipf-like distribution [6]. Similarly workloads vary from small documents (e.g., on-line book stores, browsing sites, etc.) to large documents (e.g., download sites, etc.). Further, workloads may contain requests for simple cacheable static or time invariant content or more complex dynamic or time variant content via CGI, PHP, and Java servlets with a back-end database. Due to these varying characteristics of workloads, we classify the workloads into three

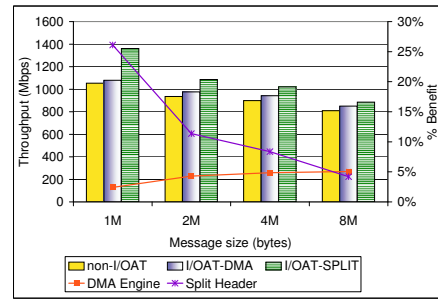



Figure 7. I/OAT split-up benefits: (a) CPU Utilization and (b) Throughput

broad categories: (i) Single-file Micro workloads, (ii) Zipf-like workloads and (iii) Dynamic content workloads. However, in this paper, we focus our analysis on the first two categories.

Single-File Micro workloads: These workloads contain only a single file. Several clients request the same file multiple times. These workloads are used to study the basic performance achievable by the data-center environment without being diluted by other interactions in more complex workloads. We use workloads ranging from 2 KB to 10 KB file sizes since this is considered the average file size for most documents in the Internet.

Zipf-like Workloads: It has been well acknowledged in the community that most workloads for data-centers hosting static content, follow a Zipf-like distribution [6]. According to Zipf law, the relative access probability of a request for the i 'th most popular document is proportional to $1/i^\alpha$, where α determines the randomness of file accesses. In our experiments, we vary α from 0.95 to 0.5, ranging from high temporal locality for documents to low temporal locality.

We also evaluate the performance achieved inside the data-center with proxy servers acting as clients and web-servers with I/OAT capability. For example, within a data-center, the proxy servers forward dynamic content requests to the application servers. The application servers are known to be CPU-intensive due to processing of scripts such as PHP, CGI, servlets, etc. If the application servers have I/OAT capability, the servers can not only accept more number of requests but can also process the pending requests at a faster rate, due to reduced CPU utilization.

For both of the scenarios, we use a testbed with one proxy at the first tier and one web-server at the second tier. Each client fires one request at a time and sends another request after getting a reply from the server.

5.2 Experimental Results

In this section, we separate our analysis into two categories. First, we analyze the performance benefits of I/OAT with the two workloads. As mentioned above, we use Testbed 2 to fire requests to the proxy server. Due to the I/OAT capability on the server nodes, we expect the performance of the servers to improve.

5.2.1 Analysis with Single File Traces

In this experiment, we use five different traces with varying file size requests. Trace 1 uses an average file size of 2 KB while

Traces 2, 3, 4 and 5 use 4 KB, 6 KB, 8 KB and 10 KB, respectively. Each client has a 10,000 request subset of different files and reports the TPS (Transactions Per Second) achieved after getting the response from the servers for all of the requests. TPS is defined as the total number of transactions serviced per second as seen by the client. Higher TPS values attribute to better server performance.

Figure 8a shows the performance of a data-center with varying trace files ranging from 2 KB to 10 KB. As shown in Figure 8a, we see that the throughput reported by I/OAT is consistently better than the throughput reported by non-I/OAT. It is also to be noted that for Trace 2, with average file size of 4 KB, we see a significant throughput improvement for I/OAT. I/OAT reports a TPS which is close to 9,754 whereas non-I/OAT reports only 8,569 TPS resulting in a 14% overall improvement. For other traces, we see around 5-8% TPS improvement with I/OAT.

5.2.2 Analysis with Zipf File Traces

Next, we evaluate the performance of the data-center with the zipf-trace. As mentioned earlier, we vary the α values from 0.95 to 0.5 ranging from high temporal locality to low temporal locality. As shown in Figure 8b, I/OAT consistently performs equal to or better than non-I/OAT. Further, we note that non-I/OAT achieves close to 1,989 TPS whereas I/OAT achieves close to 2,212 TPS, i.e., I/OAT achieves up to 11% throughput benefit as compared to non-I/OAT. Also, we measured the CPU utilization for these experiments but the improvement was negligible. However, the throughput improves since the server can accept more requests (i.e., reduction in CPU overheads result in greater CPU efficiency).

5.2.3 Analysis with Emulated Clients

Next, we evaluate the performance of I/OAT within a data-center when both the proxy and the web servers have I/OAT capability. In this experiment, the proxy server acts as a client in sending the requests to the web servers. We use only the nodes in Testbed 1. Due to reduced CPU usage in proxy nodes using I/OAT, we expect the clients to fire requests at a much faster rate, resulting in higher throughput. Figure 9 shows the performance with I/OAT capability for increasing number of client threads. In this experiment, we fix the file size to 16 KB and varied the number of client threads from 1 to 256. As shown in Figure 9, we note that the performance of I/OAT is

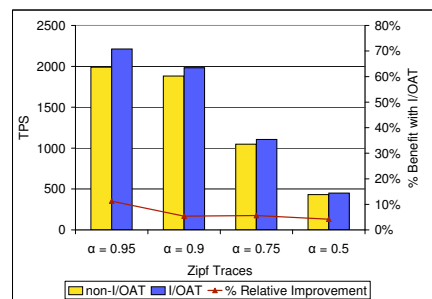
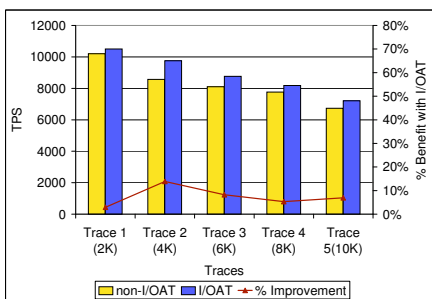


Figure 8. Performance of I/OAT and non-I/OAT: (a) Single File Traces and (b) Zipf Trace

similar to non-I/OAT from one to sixteen threads. As the number of threads firing the requests increases from 32 to 256, we observe two interesting trends. First, I/OAT throughput performance increases with an increasing number of client threads. With 256 client threads, I/OAT achieves close to 14,996 TPS whereas non-I/OAT achieves only around 12,928, i.e., I/OAT achieves close to 16% throughput improvement as compared to non-I/OAT. It is also to be noted that the CPU utilization of I/OAT is consistently lower than the CPU utilization of non-I/OAT. In this experiment, we only report the CPU utilization on the client node, since our focus is to capture the client-side benefits when clients have I/OAT capability. We observe that the CPU utilization saturates with non-I/OAT with 64 threads and thus the throughput does not improve any further with an increase in the client threads after 64. With I/OAT, we see that the CPU utilization saturates only with 256 client threads, resulting in superior performance. I/OAT not only improves the performance by 16%, but can also support a large number of threads (up to a factor of four as compared to non-I/OAT).

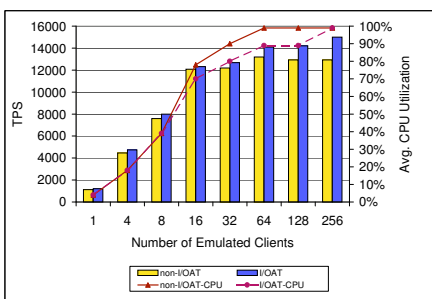


Figure 9. Clients with I/OAT capability using a 16 KB file trace

6 PVFS Performance Evaluation

In this section, we compare the PVFS performance over non-I/OAT and I/OAT with the original PVFS implementation [7]. All experiments in this section were performed on Testbed 1.

6.1 Evaluation Methodology

It is to be noted that due to the wide difference between the network and the disk speeds, the effective throughput achieved by PVFS could be limited by that of the disk. However, data used by applications is frequently in server memory, e.g., due to file caching and read-ahead mechanisms. Thus, applications can still benefit from fast networks in such cases. Hence, we

designed our experiments based on a memory-resident file system, *ramfs*. These tests are designed to stress the network data transfer independent of any disk activity.

6.2 Experimental Results

We split our evaluation into two categories. First, we design an experiment that measures the PVFS concurrent read and write bandwidth. Next, we evaluate the file system performance in the presence of multiple streams performing the bandwidth test.

6.2.1 PVFS Concurrent Read and Write on *ramfs*

The test program used for concurrent read and write performance is *pvfs-test*, which is included in the PVFS release package. We followed the same test method as described in [7]. In all tests, each compute node simultaneously reads or writes a single contiguous region of size $2N$ MB, where N is the number of I/O nodes in use. For example, if the number of I/O nodes is four, the request size is then 8 MB. Each compute node will access 2 MB of data from each I/O node. It is to be noted that since I/OAT is a receiver-side optimization, we report the average CPU utilization at the client-side while performing a read operation and report the CPU utilization at the server-side while performing the write operation.

Figure 10a shows the PVFS read bandwidth performance with I/OAT and non-I/OAT with six I/O servers. With non-I/OAT, the bandwidth increases from 361 MB/s to 649 MB/s with an increasing number of compute nodes. With I/OAT, the bandwidth increases from 360 MB/s to 731 MB/s with an increasing number of compute nodes, i.e., I/OAT achieves 12% increase in throughput as compared to non-I/OAT with six compute nodes. Further, the results are consistent with the micro-benchmark evaluations seen in Section 4. In terms of CPU utilization, we see that I/OAT achieves close to 15% improvement as compared to non-I/OAT using six compute nodes. We see similar trends when we decrease the number of I/O servers to five as shown in Figure 10b. However, the improvement seen both in terms of throughput and CPU utilization are much less than when compared to the six I/O server configuration.

Figure 11a shows the PVFS write bandwidth performance with I/OAT and non-I/OAT. Again, we see similar trends as seen with the PVFS read performance. With non-I/OAT, the write bandwidth performance increases from 464 MB/s to

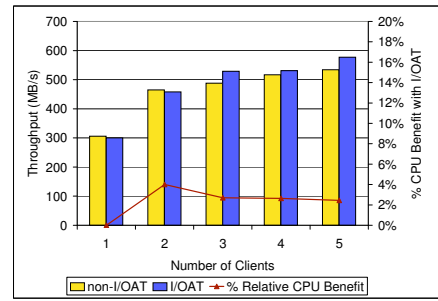
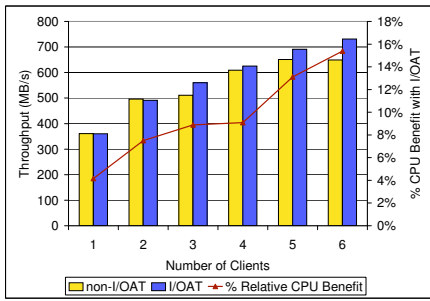


Figure 10. PVFS Concurrent Read Performance: (a) Six I/O Servers and (b) Five I/O Servers

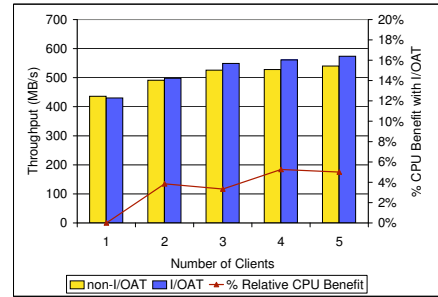



Figure 11. PVFS Concurrent Write Performance: (a) Six I/O Servers and (b) Five I/O Servers

697 MB/s. With I/OAT, the write bandwidth performance increases from 460 MB/s to 750 MB/s with an increasing number of compute nodes, i.e., I/OAT achieves 8% throughput benefit as compared to non-I/OAT. Further, I/OAT achieves close to 5% CPU benefit as compared to non-I/OAT using six compute nodes. We see similar trends when we decrease the number of I/O servers as shown in Figure 11b.

6.2.2 PVFS read performance with multiple streams

In this experiment, we emulate a scenario where multiple clients access the file system. We increase the number of clients from 1 to 64 and study the behavior of concurrent read performance in PVFS. As shown in Figure 12, we observe that the performance of I/OAT is either equal to or greater than the performance achieved by non-I/OAT. It is to be noted that we report the CPU utilization numbers measured on the client-side which also has the I/OAT capability. Due to the fact that the client with I/OAT can receive the data at a much faster rate, the clients can also fire PVFS read requests at a faster rate than compared to non-I/OAT. As a result, we observe that the CPU utilization on the client-side is around 10-12% higher than the non-I/OAT case.

In summary, PVFS with I/OAT has a 12% improvement for concurrent reads and 8% improvement for concurrent writes as compared to PVFS with non-I/OAT. In terms of CPU utilization, PVFS with I/OAT achieves 15% benefit for concurrent reads and 7% benefit for concurrent writes as compared to PVFS with non-I/OAT.

7 Discussion

As mentioned earlier, I/OAT helps in improving the network performance and reducing the CPU utilization through its features such as the asynchronous copy engine, split-headers, etc.

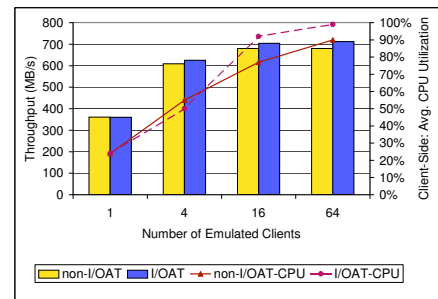


Figure 12. Multi-Stream PVFS Read Performance

The asynchronous copy engine can also be used by user applications to improve the performance of memory operations and also to improve the communication performance between two processes within the same node. Since the copy engine can operate in parallel with the host CPU, applications can achieve significant benefits by overlapping useful computations with memory operations. With the introduction of chip-level multiprocessing (CMP), per-core memory bandwidth reduces drastically, thus increasing the need for such asynchronous memory operations as a memory-latency hiding mechanism. Mohit, et. al., have proposed soft-timer techniques to reduce the receiver-side processing [5]. I/OAT can co-exist with this technology to further reduce the receiver-side overheads and improve the performance.

On the other hand, I/OAT requires special network adapters to support the split-header and the multiple receive queue feature. In addition, the host system should have an in-built copy engine to support the asynchronous copy feature. Also, the memory controller uses physical addresses and hence, the pages cannot be swapped during a copy operation. As a result, the physical pages need to be pinned before initiating the copy

operation. Due to the page-pinning requirement, the usefulness of the copy engine becomes questionable if the pinning cost exceeds the copy cost.

8 Concluding Remarks and Future Work

I/O Acceleration technology (I/OAT) developed by Intel is a set of features particularly designed to reduce the packet processing overheads on the receiver-side. This paper studies the benefits of I/OAT technology through extensive evaluations of micro-benchmarks as well as evaluations on two different application domains: (1) A multi-tier Data-Center environment and (2) A Parallel Virtual File System (PVFS). Our micro-benchmark results show that I/OAT results in 38% lower overall CPU utilization in comparison with traditional communication. Due to this reduced CPU utilization, I/OAT delivers better performance and increased network bandwidth. Our experimental results with data-centers and file systems reveal that I/OAT can improve the total number of transactions processed by 14% and throughput by 12%, respectively. In addition, I/OAT can sustain a large number of concurrent threads (up to a factor of four as compared to non-I/OAT) in data-center environments, thus increasing the scalability of the servers.

We are currently working on two broad aspects with respect to I/OAT. First, we are looking at modifying the applications to realize the true benefits of I/OAT capability. Second, we are trying to provide an asynchronous memory copy operation to user applications. Though this involves some amount of overhead (such as context switches, user page locking, etc), asynchronous memory copy can help applications in using the CPU cycles intelligently.

9 Acknowledgments

We would like to thank Intel Corporation for providing access to I/OAT-based systems and the kernel patch which includes all of their features. We would also like to thank Eric Simpson from the I/OAT team of Intel Corporation for providing us crucial insights into the implementation details.

References

- [1] Accelerating High-Speed Networking with Intel I/O Acceleration Technology. <http://www.intel.com/technology/ioacceleration/306517.pdf>.
- [2] Increasing network speeds: Technology@intel. <http://www.intel.com/technology/magazine/communications/intel-ioat-0305.htm>.
- [3] Intel I/O Acceleration Technology. <http://www.intel.com/technology/ioacceleration/306484.pdf>.
- [4] USNA, TTCP: A test of TCP and UDP Performance, August 2001.
- [5] Mohit Aron and Peter Druschel. Soft timers: efficient microsecond software timer support for network processing. *ACM Transactions on Computer Systems*, 18(3):197–228, 2000.
- [6] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker. Web caching and zipf-like distributions: Evidence and implications. In *INFOCOM (1)*, pages 126–134, 1999.
- [7] Philip H. Carns, Walter B. Ligon III, Robert B. Ross, and Rajeev Thakur. PVFS: A Parallel File System for Linux Clusters. In *Proceedings of the 4th Annual Linux Showcase and Conference*, pages 317–327, Atlanta, GA, 2000. USENIX Association.
- [8] A. Chandra, W. Gong, and P. Shenoy. Dynamic Resource Allocation for Shared Data Centers Using Online Measurements. In *Proceedings of ACM Sigmetrics 2003, San Diego, CA*, June 2003.
- [9] J. Chase, A. Gallatin, and K. Yocum. End system optimizations for high-speed TCP. *IEEE Communications Magazine*, 39(4):68–74, 2001.
- [10] Jeffrey S. Chase, Darrell C. Anderson, Prachi N. Thakar, Amin Vahdat, and Ronald P. Doyle. Managing energy and server resources in hosting centres. In *Symposium on Operating Systems Principles*, 2001.
- [11] Avery Ching, Alok Choudhary, Wei keng Liao, Robert Ross, and William Gropp. Noncontiguous I/O through PVFS. In *Proceedings of the IEEE International Conference on Cluster Computing*, 2002.
- [12] J. R. David D. Clark, Van Jacobson, and H. Salwen. An analysis of TCP processing overhead, 1989.
- [13] Annie Foong et al. TCP performance analysis revisited. In *IEEE International Symposium on Performance Analysis of Software and Systems*, March 2003.
- [14] Armando Fox, Steven D. Gribble, Yatin Chawathe, Eric A. Brewer, and Paul Gauthier. Cluster-based scalable network services. In *Symposium on Operating Systems Principles*, 1997.
- [15] Andrew Gover and Christopher Leech. Accelerating network receiver processing. <http://linux.inet.hr/files/ols2005/gover-reprint.pdf>.
- [16] S. Makineni and R. Iyer. Architectural characterization of TCP/IP packet processing on the Pentium M microprocessor. In *High Performance Computer Architecture, HPCA-10*, 2004.
- [17] G. Regnier, S. Makineni, R. Illikkal, R. Iyer, D. Minturn, R. Huggahalli, D. Newell, L. Cline, and A. Foong. TCP Onloading for Data Center Servers. In *IEEE Computer*, Nov 2004.
- [18] John Reumann, Ashish Mehra, Kang G. Shin, and Dilip Kandlur. Virtual services: A new abstraction for server consolidation. In *In Proceedings of the USENIX 2000 Technical Conference*, June 2000.
- [19] Yasushi Saito, Brian N. Bershad, and Henry M. Levy. Manageability, availability and performance in porcupine: A highly scalable, cluster-based mail service. In *Symposium on Operating Systems Principles*, 1999.
- [20] H. V. Shah, D. B. Minturn, A. Foong, G. L. McAlpine, R. S. Madukkarumukumana, and G. J. Regnier. CSP: A Novel System Architecture for Scalable Internet and Communication Services. In *the Proceedings of the 3rd USENIX Symposium on Internet Technologies and Systems*, pages pages 61–72, San Francisco, CA, March 2001.
- [21] K. Shen, H. Tang, T. Yang, and L. Chu. Integrated Resource Management for Cluster-based Internet Services. In *Proceedings of Fifth USENIX Symposium on Operating Systems Design and Implementation (OSDI '02)*, pages 225–238, Boston, MA, December 2002.
- [22] Rajeev Thakur, William Gropp, and Ewing Lusk. On Implementing MPI-IO Portably and with High Performance. In *Proceedings of the 6th Workshop on I/O in Parallel and Distributed Systems*, pages 23–32. ACM Press, May 1999.