

Characterization of E-Commerce Traffic*

Udaykiran Vallamsetty
Advanced Micro Devices
Sunnyvale, CA 94086
uday.vallamsetty@amd.com

Krishna Kant
Intel Corporation
Hillsboro, OR 97124
krishna.kant@intel.com

Prasant Mohapatra
Department of Computer Science
University of California, Davis
prasant@cs.ucdavis.edu

Abstract

The world wide web has achieved immense popularity in the business world. It is thus essential to characterize the traffic behavior at these sites, a study that will facilitate the design and development of high-performance, reliable e-commerce servers. This paper makes an effort in this direction. Aggregated traffic arriving at a Business-to-Business (B2B) and a Business-to-Consumer (B2C) e-commerce site was collected and analyzed. High degree of self-similarity was found in the traffic (higher than that observed in general web-environment). Heavy-tailed behavior of transfer times was established at both the sites. Traditionally this behavior has been attributed to the distribution of transfer sizes, which was not the case in B2C space. This implies that the heavy-tailed transfer times are actually caused by the behavior of back-end service time. In B2B space, transfer-sizes were found to be heavy-tailed. A detailed study of the traffic and load at the back-end servers was also conducted and the inferences are included in this paper.

1 Introduction

The explosive popularity of Internet has propelled its usage in several commercial avenues. E-commerce, the usage of Internet for buying and selling products, has found a major presence in today's economy. E-commerce sites provide up-to-date information and services about products to users and other businesses. Services ranging from personalized shopping to automated interaction between corporations are provided by these web-sites. It has been reported that e-commerce sites generated \$132 Billion in 2000, more than double of the \$58 Billion reported in 1999 [1]. Even though the power of the servers hosting e-commerce sites has been increasing, e-commerce sites have been unable to improve their level of service provided to the users. It has been reported that around \$420 Million has been lost in revenues due to slow processing of the transactions in 1999. Thus it is desirable and necessary to focus on the performance of the servers used in these environments.

There are two main classes of e-commerce sites, Business-to-Business (B2B) and Business-to-Consumer (B2C), providing services to corporations and individual users respectively. Web sites like Delphi, which provide services to corporations like General Motors come under B2B sites, whereas sites like Amazon.com providing services to general users come under B2C sites.

In this paper, we have analyzed the characteristics of e-commerce traffic. Traffic from a B2C and a B2B site is being used for the study. The workload is initially inspected for understanding the diurnal nature of the traffic. Different load periods were identified for both the B2C and B2B environments. These have been found to be complimentary in nature, which may be intuitive. A set of parameters were chosen for each site for each component which would impact the performance of the system to the maximum extent. Statistical tests are then used to prove the self-similar nature of the traffic at different scales. Two different tests are used for validating the results for each of the parameters. It has been observed that the arrival traffic is highly bursty in nature, much more than the burstiness seen in normal web-traffic [5]. The response-time distribution is found to be heavy-tailed. This has been previously attributed to the heavy-tailed nature of request and response file-sizes. But the behavior of transfer sizes is not heavy-tailed, unlike the general web-environment. The traffic arriving at the back-end servers is characterized to obtain similar statistics about the impact of burstiness on the system. Also preliminary tests have shown that the back-end utilization is more bursty than the front-end server utilization, the reasons for which are explained later. A correlation is drawn between the behavior of the front-end and the back-end servers under different load conditions. Performance implications from the results of the above experiments will give valuable information for improving e-commerce server performance.

The workload characterization studied in this paper is based on one representative system from each of the environments (B2C and B2B). Considering the difficulty in obtaining this valuable and guarded information from the e-commerce sites, and the fact the sites we have considered are quite busy, the results, although preliminary, could be valuable for future studies on e-commerce workload characterization and server designs.

The rest of the paper is organized as follows. Related work is outlined in Section 2. Section 3 discusses the architecture of e-commerce sites along with a description of the configuration of the sites used for this study. Section 4 discusses the behavior of the workload and the traffic and load characteristics of the front-end and back-end servers. The concluding remarks are sketched in Section 5.

Note: For this study data from two popular sites (one B2C and the other B2B) was used. Due to a non-disclosure agreement (NDA), the identity of these sites is not revealed. Throughout this work the two sites are identified as B2C site and B2B site. Without the NDA, we would not have been able to acquire the data for the study.

*This work was supported in part by the National Science Foundation.

2 Related Work

Although there have been several studies reported on the workload characterization of general web servers [3, 7, 10, 13], only a few studies have been reported on the characterization of e-commerce traffic based on the client behavior. The main reason for this shortcoming is the unavailability of representative data. E-commerce sites have highly secure information in the traces and access logs. Due to the security implications e-commerce sites are reluctant to divulge this information for research purposes. Due to this, studies in this field are still in the preliminary stages. In [8] the authors have developed a resource utilization model for a server which represents the behavior of groups of users based on their usage of the site.

It should also be noted that the existing work reported on e-commerce traffic has been done on the front-end servers only and to the best of our knowledge nothing has been reported on the back-end servers. The back-end servers are the ones which experience the maximum load in an e-commerce environment [4]. We would like to characterize the load on the back-end servers along with a study of the system characteristics collected from system logs in E-commerce sites.

3 E-Commerce Architecture

A generic organization of e-commerce sites is depicted in Figure 1. E-commerce sites can be broadly classified into two different categories. Business to Business (B2B) and Business to Consumer (B2C). The main difference between the two categories of sites lies in the user population accessing these sites. B2B sites serve transactions between different businesses whereas B2C sites serve general users over the Internet.

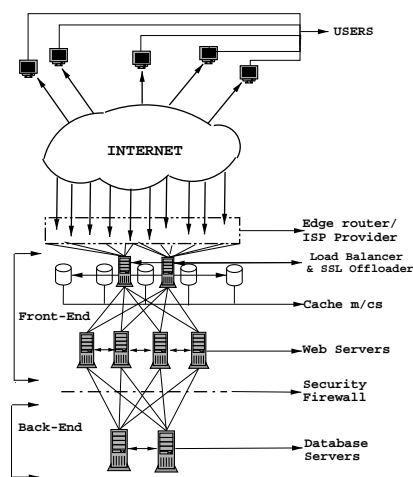


Figure 1: A Generic E-Commerce Site.

Business-to-Business: One of the main characteristics of this category of sites is the regularity in the arrival traffic. It was observed that heavy traffic comes between 9am to 5pm, normal business hours. Regularity does not imply the lack of heavy spikes in the traffic. There will be sustained

peak load on the system either due to seasonal effects or due to the availability of different services at the site. These sites can be categorized by the high amount of buying taking place in them. It has been observed that the percentage of transactions resulting in buying are very high compared to those in B2C environment.

Business-to-Consumer: B2C servers are the normal e-commerce sites where any user can get service. The security involved in B2C site is only restricted to any financial transactions involved, whereas in a B2B environment all the transactions are normally done in secure mode. One implication of this is that increased buying in a B2C environment can throttle the system since the designed system does not expect high percentage of buy transactions. Another important characteristic of a B2C site is the very low tolerance to delayed responses. This increases the need to make QoS more important than providing absolute security for all the transactions, hence security is reserved for transactions involving buying.

3.1 Front-End and Back-End Servers

Typically the front-end servers are comprised of the web server, application server, server load balancer, and the secure socket layer (SSL) off-loader. Front-end web servers serve requests from the clients and are the only authorized hosts able to access the back-end database and application services as necessary. The application servers are responsible for the business logic services. The application server will be the most heavily loaded server in the B2C environment. This is due to the heavy traffic of dynamic and secure requests arriving at the server. In a large scale e-commerce site, there will be dedicated application servers, alternatively these servers can be combined with the Web Servers or the Database servers. Due to the heavy traffic seen by e-commerce servers and also due to the availability requirements, there will be a network of web servers instead of a single monolithic server at the front-end. This basically improves the scalability and fault-tolerance of the server to any bursts of busy traffic. Load balancers help increase the scalability of an e-commerce site. Load balancing works by distributing user requests among a group of servers that appear as single virtual server to the end user. SSL is a user authentication protocol developed by Netscape using RSA Data Security's encryption technology. Many commerce transaction-oriented web sites that request credit card or personal information use SSL. The SSL off-loader typically decrypts all *https* requests arriving at the server.

The back-end servers mainly comprise of the database servers and the firewall which would protect sensitive data from being accessed by unauthorized clients. These firewalls provide security services through connection control. They are predominantly used when protecting mission-critical or sensitive data is of utmost importance. The database servers reside in the back-end of the network and house the data for e-commerce transactions as well as sensitive customer information. This is commonly referred to as the data services. The clients do not directly connect to these servers, the front-end Web servers initiate connections to these servers when a client conducts a series of actions

such as logging in, checking inventory, or placing an order. Most e-commerce sites scale up their database servers for scalability and implement fail-over clustering for high availability. Partitioned databases, where segments of data are stored on separate database servers, are also used to enhance scalability and high availability in a scale-out fashion.

3.2 B2C Configuration

A simplified configuration of the B2C site being used for the study is given in Figure 2. This site comprises of ten web servers, each one powered by a Intel Quad P-III systems with a 512MB of RAM. The web servers run IIS 4.0 HTTP server. This cluster of web servers is supported by three image servers, each one powered by a Dual P-II system. As can be seen from the figure, the image servers serve both the database servers and the front-end web servers. For the purpose of our study, the image servers were considered to be in the back-end system. The product catalog server, connected to both the front-end and the back-end, runs an NT 4.0 providing backup and SMTP services to the back-end servers. The LDAP server is connected to the back-end.

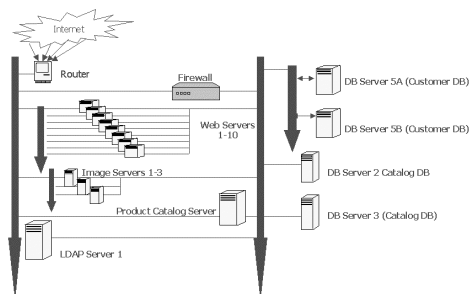


Figure 2: Simplified configuration of the B2C site.

3.3 B2B Configuration

In the B2B space, the design of e-commerce sites is completely different from their design in B2C space. Here the user population is known a-priori. The transactions being processed by each user arriving at the server is also known with reasonable bounds. B2B sites serve a limited population as opposed to B2C sites which aim at serving the entire Internet. These aspects enable the designers to customize the site to specific user requirements.

Scalability is one of the main issue that has to be taken care of when designing such customized system. So the design is done as a cluster of B2C sites, interconnected to form a large B2B portal. The interconnections between the individual B2C components in the site determine the user population to that site and also the services provided by that site. Figure 3 shows a simplified version of the B2B site being used for the study. Each of the web servers, can be individually used as a B2C site with its own database and network connection.

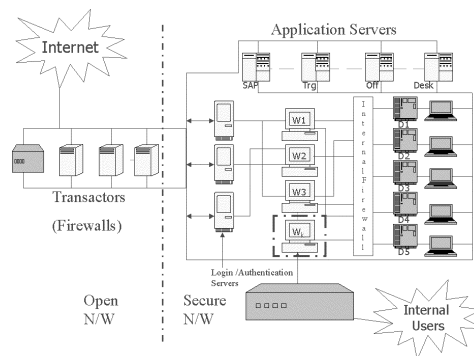


Figure 3: Simplified configuration of the B2B site.

4 Workload Characterization

In this study we have analyzed the behavior of e-commerce servers with relation to the behavior of the incoming traffic. Data was collected at different levels in the system. Web Server access logs from the the front-end and the back-end servers were collected at a granularity of 1 sec. This is an application level data giving the load on the httpd. This data will give the characteristics of the traffic arriving at the system, average network bandwidth utilization, and the file transfer rate.

For the system level information, data was collected from the Performance logs from all the servers present in the site. This data was collected at a granularity of 5 sec. This would give information about the I/O bandwidth used, the processor and disk utilization of the system etc.. Data was collected at a constant rate of 5 sec intervals. So this data is at a higher scale than the logs from the web servers. But both the scales are below the non-stationarity time scale used for the analysis.

Data was collected at the server and the performance monitor for an entire day. A weekday is used for data collection since this would represent the average traffic. Additionally, data for a five day period was used to study the average behavior of the traffic over a long period of time.

4.1 Characteristics of the Workload

The main differences between general web and e-commerce workload are the following.

1. Presence of a high level of Online Transaction Processing (OLTP) activity is observed among the transactions at the server. This is due to the database transactions accruing for every request from the user. Due to security reasons most of the data is present in the database server which is protected by a secure firewall. This prevents the web server from responding to most of the requests without sending a query to the back-end server.
2. A large proportion of requests come in secure mode. B2C traffic has lesser secure traffic, B2B sites experience almost complete secure traffic from users. This is due to the heavy security constraints present in industry to industry transactions. Increased amount of

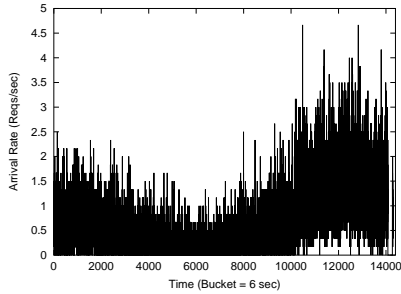


Figure 4: Arrival Process at B2C site.

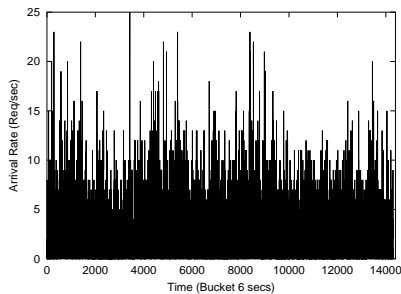


Figure 5: Arrival process at B2B site.

secure transactions implies heavy processing load at the front-end server. Most of the sites have SSL offloaders, which do encryption/decryption of requests to reduce the load on the system. This process adds to the response time. Aggregating these transactions with normal transactions increases the variability in the response times observed by the user.

3. The proportion of dynamic requests (that require some amount of processing) is very high, as was expected. In fact, in most e-commerce sites almost all requests are handled as dynamic requests.

4.2 Front-End Characterization

A visual inspection reveals the workload at e-commerce sites to be more bursty than normal web workload. To study this behavior, the following parameters were used, which would have the maximum impact on the behavior of the traffic: Arrival process, Utilization of the server, Response time, Request file sizes, and Response file sizes.

4.2.1 Arrival Process

Figures 4 and 5 show the arrival process at the B2C and B2B e-commerce sites for an entire day (12am-12am). The data shows traffic on a normal weekday with an average arrival rate of 0.65 requests/sec at the front-end web server for the B2C site and around 1 req/sec arrival rate at one of the web servers in the B2B site. A visual inspection reveals the burstiness in the arrival process. The B2C server is a 4P system with an average processor utilization of 6% per

processor and disk utilization of 2% during the period starting from 9.00am till 6.00pm. The low utilization is typical of e-commerce sites since they are designed for much higher load and sustain a very minimal load during normal working periods. It is the high load periods showing bursts of orders of magnitude more than normal operating parameters which cause concern for better capacity planning and performance analysis of these systems.

Figures 4 and 5 show that the sites have distinct high and low load periods during the course of a day. For the B2C site, busy period starts around 6:00pm in the evening and ends at around 11:00pm in the night. Since this is a B2C site serving general consumers, the traffic is heavy during the after-office periods. Distinctive low periods during the morning between 7:30am to 11:30am can also be observed. In case of the B2B site, the traffic concentration lies mostly during normal office hours, between 9:00am to 8:00pm, which is intuitive. It should be noticed that the graphs show aggregated arrival traffic for the B2B site and the averaged arrival process for the B2C site.

The Arby-Veitch (AV) [11] estimator test was used for estimating the Hurst-parameter (H-parameter) [6] for the arrival time-series. This is known to be a reliable test for workloads with busy periods showing a non-stationary behavior. Hurst parameter is also calculated using the R/S plot test [6]. Reliability of this test under low time-scales for e-commerce traffic is tested by comparing the H-parameters obtained using the two methods.

The H-parameter is estimated to be 0.662 using the AV test. This shows that the arrival process at the B2C site is *self-similar* in nature. The Hurst parameter is also estimated to be 0.662 using the R/S plot test for the B2C site, which matches the estimation made by the AV-estimator. Similar test was done for the arrival traffic at the B2B site. Using the AV-estimator the H-parameter was estimated at 0.69, whereas the R/S plot gave an estimate of 0.70 for the H-parameter, which is a good approximation.

4.2.2 Processor Utilization

Figures 6 and 7 show the utilization of the front-end web server for the B2C and B2B sites respectively. As explained earlier, this data is collected between 9:00am till 5:00pm at a granularity of 5 secs for the B2C site. For the B2B site the data represents the activity between 10:00 am in the morning till 9:30 am the next day morning. The B2C server sustains a constant load throughout the day, with an average load of 7% on each of the four processors. High and low load periods can be observed on the B2C server during the course of the day. This behavior is absent in the B2B server. This is due to the a-priori knowledge of the transactions and load from users in the B2B space. B2B sites are customized for specific traffic patterns and a normal traffic would not affect the load on the system to a higher degree. Thus the load on the system appears almost constant even though there is a variation in the arrival rate at the server. The time-series obtained from the utilization was also tested for self-similar behavior. The AV-wavelet based test and the R/S plot test are used for estimating the H-parameter. The estimated H-parameter is 0.755 using the AV estimator, and 0.77 using

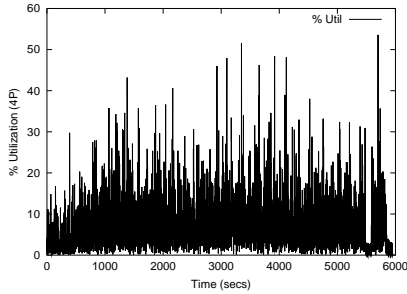


Figure 6: Processor(4P) Utilization of the front-end web-server at B2C site

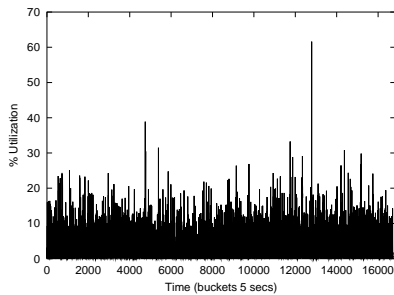


Figure 7: Processor(2P) Utilization of one of the front-end web-servers at B2B site

the R/S plot test for the B2C site. In the B2B space, the load on the system did not have a high degree of self-similarity. The H-parameter is estimated to be 0.66 using both the AV-estimator and the R/S plot test. Due to a balanced load on the B2B system throughout the duration, the degree of self-similarity is very low. The effect of the arrival process is not seen in the overall load sustained by the B2B server.

A higher H-parameter implies an increased degree of self-similarity. Utilization is a factor of the response-time and the arrival process. The inherent burstiness in the arrival process is already established in the previous section. The higher H value can be attributed to long-range dependence in the service process.

4.2.3 Response Time

In Figure 8, the response time observed by the users over the entire day period is shown for the B2C site. Previous studies [5, 12] have concentrated on the study of the heavy-tailed behavior of web response times. In this work the response-time distribution is converted into a time-series by aggregating the response-times seen for non-overlapping intervals of 5 secs. Even though the times seen are not the actual response times observed by the user, they can be used for time-series analysis. Only a multiplicative factor of 1/5 will be required to get the actual response-times. The time-series obtained is checked for self-similarity and any non-stationary behavior. The AV test and R/S plot test are used for estimating the H-parameter (A-V test $H = 0.545$, R/S Plot $H = 0.560$). As explained earlier, a good estimation of

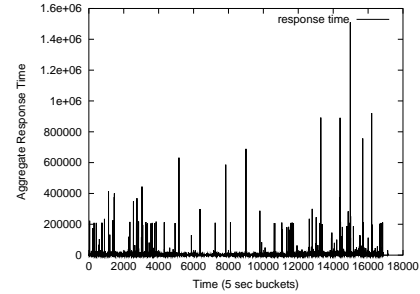


Figure 8: Aggregate Response time at the front-end Web-Server (4P), B2C site

H-parameter is obtained using R/S test, only when the time-series is stationary. So both the tests are used for estimating the H-parameter.

Response time is one of the very important performance metrics in the design and analysis of any server system. High burstiness in the arrival traffic implies saturating server queues, leading to high response times. Studies have shown that the 90th percentile response-times can be used for predicting the mean response-time [2]. This measure cannot be used in presence of high burstiness in the response-time distribution. Figure 8 shows response times orders of magnitude higher during the high load periods in the evening. Comparing this graph with the arrival process shown in Figure 4, unmistakable correlation can be found between the different load periods. Even though the utilization of the system does not get effected, buffer queue lengths increase thereby increasing the user perceived response times. Increased burstiness impacts the overall response time of the system to a higher extent than the arrival process. This burstiness in the response time is a factor of the back-end data retrieval time and the server processing time.

4.2.4 Request/Response File Sizes

The request and response file sizes in web environment [5, 9] have been studied previously. It was observed that these distributions show a heavy-tailed behavior with a tail weight of approximately $\alpha = 1.06$ for file-sizes greater than 1000 bytes [5]. This was considered one of the main reasons for the heavy-tailed behavior of the web response times. In e-commerce environment, it has already been shown that transfer times have a heavy tailed behavior with $\alpha = 1.55$. In this section the behavior of transfer size distribution is studied. Figures 9 and 10 show the request and response size distribution over the observation period at the B2C server.

It can be observed that the distribution of transfer sizes is fairly constant in the B2C environment. A visual inspection rules out the possibility of heavy burstiness in the aggregated time-series obtained from the transfer sizes. The distribution of request sizes is further investigated for heavy-tailed behavior using Log-Log cumulative distribution plots (LLCD) plots [5]. Figure 11 shows the log-scale plot of the cumulative probability function over the different request sizes observed. The plot appears linear after $x > 2.5$. A linear-regression fit to the points for requests more than 320 Bytes

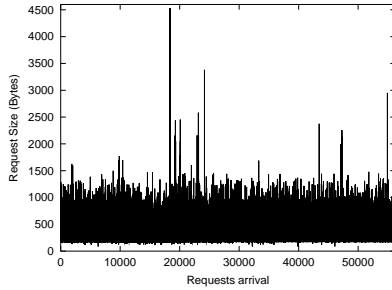


Figure 9: Request size distribution over time, B2C site

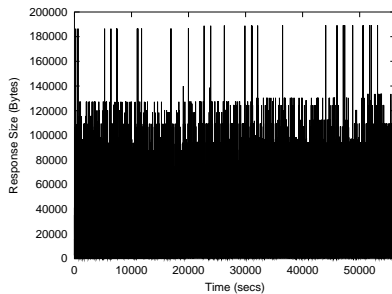


Figure 10: Response size distribution over time, B2C site

gives a line with slope $\alpha = -4.12$ ($R^2 = 0.947$). This gives an estimate of $\alpha = 4.12$ thereby indicating that the request size distribution is not heavy-tailed in nature. This result refutes the previous results about web traffic. In [5] the authors found that the requests also follow a heavy tail distribution with $\alpha = 1.16$. Using similar tests, we also infer that the response file sizes do not follow a heavy-tailed distribution.

4.3 Performance Implications

Previous studies on web traffic and LAN traffic have attributed the self-similar behavior of network traffic to the aggregation of long-range dependent ON/OFF processes. In e-commerce space, the response-times are found to be heavy-tailed in nature even though the request and response file sizes are almost a constant. The heavy-tailed behavior of response-times in web environment was believed to be caused by the heavy-tailed behavior of the file transfer sizes in the web environment. In e-commerce environment, the transfer sizes do not follow a heavy-tailed distribution as shown earlier in this section. Heavy-tailed behavior of web transfer sizes are fundamentally caused by the inclusion of image and video files in the overall traffic. Since these files are minimized in e-commerce environment (for reducing the overhead in response times), the behavior of the transfer sizes becomes somewhat intuitive. The lack of large image and video files removes the heavy-tailed nature of e-commerce traffic.

It is observed that the response time still shows a heavy-tailed behavior in both B2C and B2B space. As explained earlier this implies that the user perceived response-time can increase by orders of magnitude under load conditions. Due

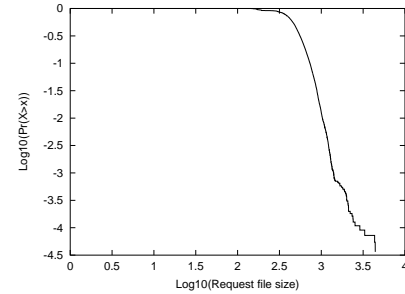


Figure 11: LLCD of request size distribution, B2C site

to the critical nature of e-commerce applications and also the business model (increasing criticality with the increase in load), it is imperative that the response-times are kept under normal bounds even in high load conditions. In e-commerce environment response-time is mainly dependent on the processing time and the transfer time. Since the file-sizes do not follow a heavy-tailed distribution, it can be safely assumed that the transfer time does not contribute to the variation in the response-time. This shows that the characteristic of the processing time is affecting the response-time to a higher extent than the response size. Also the effect of file-sizes appears to be negligible on the end-end response-times observed. This result contradicts the behavior of response-times for normal web traffic where the response-size of files can be assumed as a good approximation of the response-time. The difference is that, in web environment the transfer times consumes most portion of the response-time which is not the case in e-commerce environment due to the different composition of requests.

4.4 Back-End Characterization

The most important and sensitive information in e-commerce servers is kept in the back-end servers. It is the back-end servers that execute the business logic for the e-commerce site and are hence the most crucial components of any e-commerce server. The parameters used for doing the characterization depend mostly on the configuration of the site and the purpose of the individual components [13] in the back-end. The composition of back-end servers is closely dictated by the business model of the site. So different parameters might be interesting for different sites. In this study processor utilization and disk accesses are used for studying the characteristics of the two sites.

In the B2C site there are four different servers at the back-end. These are: Main database server, Customer database server, Image server, and LDAP server. The image server and LDAP server are not heavily loaded during the observation period. There is a single burst of traffic to and from these servers when the data is updated daily. This burst is also seen in other back-end databases and will be discussed in detail later in this chapter. The only servers that experience a sustained load throughout the day are the customer database and the main database. These two servers are used for studying the characteristics of the back-end system.

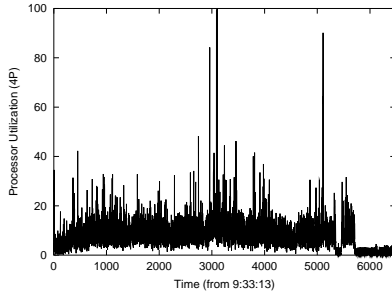


Figure 12: Processor utilization of Catalog Server (5 secs), B2B site

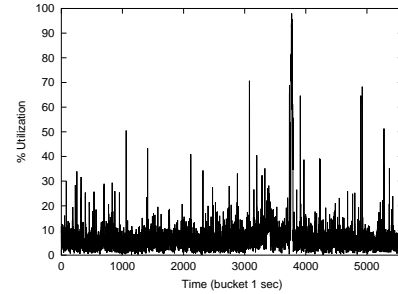


Figure 14: Processor Utilization of the B2B Database server

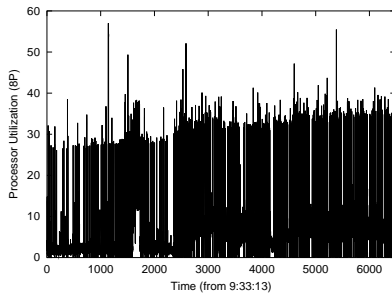


Figure 13: Processor utilization of the Main D/B server (5 secs), B2B site

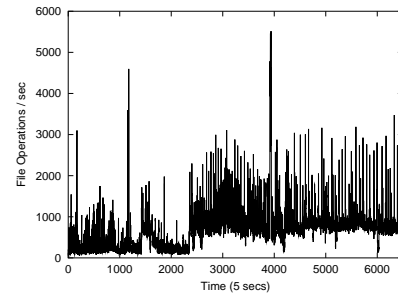


Figure 15: File Operations per second from Main DB server (5sec)

4.4.1 Processor Utilization

In Figures 12 and 13 the processor utilization of the two back-end servers in the B2C site is shown. It can be observed that the back-end server experiences a sustained load of 10% on average over the entire period. There is a visible peak of almost 100% utilization of the catalog server. This will be discussed later in the section. For the Main D/B server, the utilization remains at around 30% for most of the observation period. This shows that the load on back-end servers is higher than on the front-end servers, when compared with figure 6. Previous studies have speculated that the load on the back-end servers is more regulated due to the presence of the front-end server. One of the reasons for this speculation is the service time of the front-end server. This either causes a delay or reduces peak of any burst reaching the back-end servers. This behavior of the back-end servers is investigated by looking at the time-series obtained from the utilization of the servers.

H-parameter values of 0.87 and 0.77 were obtained for the utilization of the main database server and the catalog server respectively. The burstiness observed at the back-end servers is more than the front-end servers ($H = 0.77$). Similar results have been observed in the B2B space also. The utilization of the database server of the B2B site is shown in Figure 14. It can be observed that the load on the system reaches 100% around the 4000th bucket. This is due to the update activity which takes place periodically in most e-commerce sites. The actual time when this takes place is around 1.00 pm in the night. Similar activity can be seen in the other back-end servers. Nothing can be observed at

the front-end servers, as the bulk of the data which needs any maintenance is present in the back-end servers only. The back-end server in B2B space is also found to be more bursty than the front-end server. This contradicts previous assumptions about burstiness at the back-end servers in web environment.

4.4.2 Disk Accesses

The B2C site has four disks for the Main D/B system. Disk accesses are used for the study instead of disk utilization. Reliable data could not be obtained for the disk utilization due to the presence of a cluster of four disks.

Figures 15, shows the distribution of the file request rate at the Main D/B server. This shows the arrival rate of file requests seen by the four hard disks. Figure 16 shows the

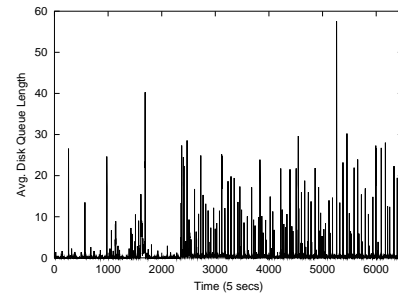


Figure 16: Disk Queue Length at Main DB server (5sec)

average queue length seen by the hard disks at the Main D/B server. The average queue length is found to be self-similar in nature with $H = 0.77$. This would result in a heavy-tailed behavior in the average response-time of the hard disk. The reason for the burstiness in the queue length can be attributed to the arrival of file transfers at the hard disk. This rate is also found to be bursty in nature with $H = 0.83$. The buffer cache does not appear to be effective since the hard disk is experiencing requests at this level of burstiness.

In the previous section, the response-time at the front-end is found to be heavy-tailed in nature even though the request and response size did not follow this distribution. The burstiness in the service time at the back-end was attributed to this behavior. Here it can be seen that the heavy-tailed distribution of response-time at the back-end is due to the bursty arrival process to the hard disks, causing the queue length to be bursty. This high burstiness in queue length will remove the effect file sizes may have on the transfer times. This conclusion also supports the previous speculation that file-sizes were not a good representation of response-times in e-commerce environment.

5 Conclusion

Aggregated traffic arriving at B2B and B2C e-commerce servers is characterized in this paper. Access logs from the web servers is collected for application level information, Microsoft performance logs were collected for system level information and processor counters were collected for architectural information. Information from this data was used to understand the load behavior of the traffic for a normal weekday. Only a specific set of parameters (arrival process, utilization, response-time, transfer sizes etc.) which would impact the system to the maximum extent were used for characterization of the workload.

Self-similar nature of the traffic was established using Hurst-parameter as a measure of degree of self-similarity. Two different tests were used for measuring the Hurst-parameter: AV-estimator and the R/S plot. It was observed that the load behavior of the two sites was complementary in nature with traffic load shifting from one type of e-commerce site to the other during the later part of the day. Unlike previous speculation, the back-end server was found more bursty than the front-end server, this was attributed to the fractal nature of the service time at the back-end. At both the sites, the response-times were found to be heavy-tailed in nature, complying to the results found in web environment. But in the B2C environment, highly bursty arrival of file requests was seen at the disks. It was found that this arrival process is causing high queuing delays at the disk reducing the impact of disk transfer time as compared to the queuing time. This increased the burstiness in the overall response-time seen at the front-end server.

This work provides an understanding of the complexity of the traffic arriving at e-commerce sites while providing a preliminary workload characterization.

References

- [1] "Real numbers behind 'net profits." ActivMedia, June 2000.
- [2] D. Krishnamurthy and J. Rolia, "Predicting the performance of an e-commerce server: Those mean percentiles," in *Proc. First Workshop on Internet Server Performance, ACM SIGMETRICS '98*, June 1998.
- [3] X. Chen, H. Chen, and P. Mohapatra, "An admission control scheme for predictable server response time for web accesses," in *Proceedings of the 10th World Wide Web Conference*, 2001.
- [4] "Cisco and microsoft e-commerce framework architecture." <http://www.microsoft.com/technet/ecommerce/ciscomef.asp>.
- [5] M. Crovella and A. Bestavros, "Self-similarity in world-wide traffic: Evidence and possible causes," *IEEE/ACM Transactions on Networking*, vol. 5, pp. 835–846, December 1997.
- [6] K. Kant, *Introduction to computer system performance evaluation*. New York, NY: Mc Graw-Hill, Inc, 1992.
- [7] K. Kant and Y. Won, "Server capacity planning for web traffic workload," *IEEE trans. on knowledge and data engineering*, pp. 731 – 747, October 1999.
- [8] D. A. Menasce, A. A. F. Almeida, R. Fonseca, and M. A. Mendes, "Resource management policies for e-commerce servers," *2nd Workshop on Internet Server Performance*, May 1999.
- [9] M. F. Arlitt and C. L. Williamson, "Web server workload characterization: The search for invariants," *ACM SIGMETRICS Conf.*, May 1996.
- [10] P. Bradford and M. Crovella, "Generating representative web workloads for network and server performance evaluation," in *Proc. 1998 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, Madison*, July 1998.
- [11] D. Veitch and P. Abry, "A wavelet based joint estimator for the parameters of lrd," *Special issue on Multiscale Statistical Signal Analysis and its Applications, IEEE Trans. Info. Th.*, vol. 45, April 1999.
- [12] Z. Sahinoglu and S. Tekinay, "On multimedia networks: Self-similar traffic and network performance," *IEEE Communications Magazine*, January 1999.
- [13] D. A. Menasce and A. A. F. Almeida, *Capacity Planning for Web Performance: Metrics, Models and Methods*. Prentice Hall, 1998.