CoSAC: Coordinated Session-based Admission Control for Multi-tier Internet Applications

Sireesha Muppala and Xiaobo Zhou

Department of Computer Science

University of Colorado at Colorado Springs, CO 80918, USA

{sdasaraj, xzhou}@uccs.edu

Abstract—Popular Internet applications deploy a multi-tier architecture, with each tier provisioning a certain functionality to its preceding tier. The session-based admission control approach (SBAC) designed for a single Web server is not effective for a multi-tier architecture. This is due to the fact that the bottleneck in a multi-tier website dynamically shifts among tiers as client access patterns change. Admission control based on only the bottleneck tier is not efficient as different sessions impose different resource consumptions at the different tiers. In this paper, we propose a multi-tier measurement based admission control (MBAC), which pro-actively accepts different session mixes based on the utilization state of all tiers. More significantly, we design a novel coordinated session-based admission control approach (CoSAC) based on a machine learning technique. It uses a Bayesian network to correlate the states of all tiers. The probability with which a session is admitted is determined by the probabilistic inference of the network after applying the evidence in terms of utilization and processing time at each tier to the network. We compare CoSAC with MBAC and a black-box approach tailored from SBAC, using the industry standard TPC-W benchmark in a typical three-tier e-commerce website. Experimental results demonstrate the superior performance of CoSAC with respect to the effective session throughput.

I. INTRODUCTION

Today, popular Internet applications deploy a multi-tier architecture, with each tier provisioning a certain functionality to its preceding tier and making use of the functionality provided by its successor to carry out its part of the overall request processing [1], [4], [7], [8], [12], [13], [14], [16]. The multi-tier computer systems keep growing in scale and complexity. They become so complicated that it is even a big challenge to get a good understanding of the entire system dynamic behaviors [11].

Admission control is a critical mechanism for peak load management and quality-of-service provisioning for Internet services. Most Internet applications are session-based. A session is a sequence of individual requests of different types made by a customer during a single visit to a website. SBAC (session-based admission control) is an important approach on e-commerce websites [2]. It proposes to use session throughput, instead of request throughput, to evaluate the web server performance. Now session throughput is widely used as a major performance metric of web applications from a service provider's perspective. However, SBAC is not effective in a multi-tier architecture. This is mainly due to the fact that the bottleneck tier in a multi-tier website dynamically shifts among tiers as client access pattern changes.

The industry TPC-W benchmark has browsing, shopping, and ordering mixes. One may argue that under heavy load conditions, only ordering sessions be accepted since they are more likely than other sessions to result in economic benefits for an e-commerce site. But our experiments found that accepting only ordering sessions under a heavy load condition is not the right way to utilize the resources in a multi-tier website. Processing a request involves multiple system components at the different tiers. Requests of different session mixes impose different resource consumptions at the different tiers. For example, studies found that the browsing related requests tend to put more pressure on the backend database server while the ordering related requests tend to put the least pressure on the database server [11]. Therefore, admission control simply based on the bottleneck tier is not efficient.

In this paper, we design and evaluate two new session-based admission control approaches for multi-tier Internet applications. We first propose a multi-tier measurement based admission control (MBAC). It proportionally accepts different session mixes based on the utilization state of the tiers. The motivation is that the saturation of the system in processing one type of requests may not necessarily mean that the system cannot handle other types of requests. Accepting a mixture of different types of sessions may make the multi-tier servers more balanced. Thus, the effective session throughput defined as the number of completed sessions can be improved.

More significantly, we propose a novel coordinated session-based admission control approach (CoSAC) with a machine learning technique. As the multi-tier server system becomes more and more complicated, people envision that empirical models built using statistical learning have great potentials to help overcome the challenges of scale and complexity [1], [11], [15]. CoSAC achieves the admission control coordination among all tiers by modeling the multi-tier system as a Bayesian network. The probability with which a session is admitted is determined by the probabilistic inference of the network after applying the evidence in terms of utilization and processing time at each tier to the Bayesian network. This results in coordinated admission decisions, by taking into account the state of all tiers at the time of admission.

We compare CoSAC with MBAC and a black-box approach tailored from SBAC. We evaluate the approaches using the TPC-W workload benchmark in a typical three-tier e-commerce environment. Experimental results demonstrate the

superior performance of CoSAC. It can improve the effective session throughput by about 50% compared to the black-box approach in most scenarios, while MBAC can improve that effective session throughput by about 20%.

II. RELATED WORK

Multi-tier server architectures impose challenges including shifting bottleneck tier, inter-tier dependency, and per-tier concurrency limit. Resource management for multi-tier Internet applications has been a very active research topic [7], [8], [12], [13], [14]. There are studies that focused on the modeling and analysis of multi-tier servers [4], [12] and dynamic server provisioning [9], [13], [14] with queueing foundations. Our work focuses on the session-based admission control for multi-tier applications, which complements the above studies.

SBAC is an innovative work on session-based admission control on e-commerce websites [2]. Simulation shows that an overloaded web server can experience a severe loss of throughput measured as a number of completed sessions compared against the server throughput measured in completed requests. SBAC is able to provide a fair guarantee of session completion, for any accepted session, independent of a session length. However, it was designed for overload control in a single web server. It is not effective in a multi-tier server system as the bottleneck tier shifts among tiers when access patterns change dynamically.

The work in [5] deploys an admission controller between the application tier and the bottleneck database tier in a three-tier website. It identifies different types of servlets and performs overload protection and preferential request scheduling in the form of shortest job first. The work in [17] studied a load shedding mechanism for busy Internet services. It deploys a selective early request termination mechanism to actively detect and abort overdue long requests to improve system throughput. The work in [6] recognizes the fact that the cost of establishing a new Secure Socket Layer (SSL) connection is much greater than that of a resumed SSL connection. The new approach prioritizes the resumed SSL connection for performance improvement. The mechanisms discussed above however do not provide any coordination between multiple tiers for session-based admission control.

Statistical machine learning techniques are recently used to measure the capacity of multi-tier Websites. One approach in [11] uses a Bayesian network to correlate low level instrumentation data such as system and user CPU time, available memory size, and I/O status that are collected at run-time to high level system states in each tier of a multi-tier web site. A decision tree is induced to identify the bottleneck dynamically when the system is overloaded. The work in [1] applies the K-nearest-neighbors machine learning approach for adding database replicas in dynamic content Web server clusters. Those studies have demonstrated the effectiveness of using machine learning techniques in system performance detection. In this paper, we use a machine learning technique for coordinated admission control on a multi-tier architecture to improve the effective session throughput.

TABLE I
REQUEST COMPOSITIONS IN TPC-W.

	Browsing	Shopping	Ordering
browsing request	95%	80%	50%
ordering request	5%	20%	50%

Fig. 1. Dynamics of a three-tier website.

III. SESSION ADMISSION CONTROL

A. Dynamics of a multi-tier website

A typical e-commerce application consists of three tiers; a front-end Web tier that is responsible for HTTP request processing, a middle application tier that implements core application functionality say based on Java Enterprise platform, and a backend database that stores product catalogs and user orders. In this context, an incoming user request undergoes HTTP processing, application server processing, and triggers queries or transactions at the database.

We simulated the activities of an e-commerce website using TPC-W workloads to demonstrate that the bottleneck tier of the three-tier website dynamically shifts. Table I shows the TPC-W workload that is composed of three distinct session mixes, Browsing, Shopping and Ordering. Each of these mixes is characterized by different probability based session navigational patterns. Sessions belonging to different mixes visit the tiers in varying number of times and pose different workloads, thereby leading to imbalance in the tier utilizations.

Figure 1 depicts the utilizations measured at different sampling intervals. A server is considered to be the bottleneck server if its utilization exceeds a pre-configured threshold. As shown by Table II, a different tier becomes the bottleneck tier at certain intervals. For example, during the sampling interval 9, the application tier is the bottleneck, while during the interval 41, the database tier is the bottleneck. The experimental results demonstrate the dynamics challenge for effective session-based admission control on a multi-tier architecture.

TABLE II
THE SHIFTING BOTTLENECK TIER.

Intervals	Bottleneck Tier
41	Database
9	Web

B. A Blackbox Approach

This approach is a straightforward extension of SBAC [2]. SBAC was designed for a single web server. This blackbox approach extends it to a multi-tier website. The utilization of various tiers is constantly monitored at sampling intervals. Using the measured utilizations in the recent past intervals, the predicted utilization for each tier for the next interval is calculated. The predicted utilizations are compared with the pre-configured tier-specific utilization thresholds. Based on the comparison results, the admission control decision is made whether to accept or reject new sessions in the next interval. In this approach, as soon as the predicted utilization of one tier exceeds the maximum threshold, it stops accepting new sessions until the predicted utilization falls below that threshold in subsequent intervals. Essentially, the admission control treats the multi-tier website as a blackbox. The admission decision is based on the utilization of the bottleneck tier, whichever it is in a multi-tier website.

C. The MBAC Approach

In a multi-tier e-commerce website, processing a request involves multiple system component in different tiers. Saturation of the system in the processing of one type of requests may not necessarily mean it cannot handle other requests. With the Blackbox approach, in any given interval, either all of the new sessions are accepted or none of them are accepted. This can lead to under utilization of the system resources at certain tiers, while only one of the tiers is overloaded and the other tiers are operating at a normal load. We propose a multi-tier measurement based admission control (MBAC) approach. It aims to overcome the limitation by accepting different traffic mixes based on the utilizations of the individual tiers.

The idea is to pro-actively accept different session mixes in proportion to the predicted utilizations at the different tiers. This is motivated by the observation that the requests of one mix type put more pressure on a specific tier while the requests of another mix type tend to put the least pressure on that tier [11]. Based on the statistical analysis of workload characteristics, MBAC may achieve more balanced utilization of the tier resources and improve the effective session throughput (completed sessions) by accepting a mixture of different sessions. Two utilization thresholds, one for minimum and one for maximum are maintained. Once the minimum utilization of one tier is exceeded, incoming sessions belonging to different traffic mixes are proportionally accepted, to keep the tier utilizations in balance and improve session throughput. For more technical details, please refer to the technical report [10].

D. The CoSAC approach

The Blackbox approach and the MBAC approach suffer from the following issues:

- The admission decision is completely based on measured values of the utilizations at individual tiers. There is no coordination between the states of multiple tiers.
- The admission control decision is executed at the interval edges. At the time of the admission control decision, if

Fig. 2. The system architecture.

Fig. 3. A Bayesian network and an instance.

it is decided that no new session will be accepted, none will be accepted until the decision is changed at end of the interval, even if some tiers are no longer overloaded during the interval.

We design the CoSAC approach with a machine learning technique to overcome the above shortcomings. Figure 2 illustrates the system architecture. Instead of making the admission control decision at a specific interval, the decision will be made each time a new session enters the system. Coordination between various tiers is achieved by modeling the system as a Bayesian network. The probability with which a session is accepted or rejected is determined by the probabilistic inference of the network after applying the evidence in terms of utilization and processing time at each tier to the network. This leads to a coordinated admission decision, by taking into account the state of all three tiers at the time of decision.

A Bayesian network is a probabilistic graphical model that represents a set of variables as nodes and their conditional probabilistic dependencies as arcs between them, in a parent-child hierarchy. The quantitative relationships between the parent and child nodes are captured by means of conditional probability tables (CPT). The CPT of each child node, captures a collection of probability distributions over the child node, one for each different parental configuration, thus quantifying the parent-child dependency.

The Bayesian network representation of the multi-tier system is illustrated in Figure 3. Each oval represents a parameter of the system, with the top portion specifying the name and the bottom portion specifying mutually exclusive and exhaustive valid states of the Bayesian network node. Each tier has a utilization parameter and a processing time parameter, that lead to a more comprehensive representation of the workload states of the tier. For example, the utilization parameter measured at the web tier is represented by the network node WebTierUtilization. It can be in either Below Threshold (BT), within Normal Range (NR) or Above Threshold (AT) state. The processing time parameter at the web site is presented the network node WebTierProcessingTime. It can also be in either BT, NR or AT states.

The numerical values shown in Figure 3 are the probability of the network node being in the state when a specific evidence in terms of a measured utilization or processing time is applied to the Bayesian network. For instance, when one new session arrives, the utilization and processing time parameters of each tier are measured, which are shown in Table III. The measured value of the web tier utilization is applied as an evidence for WebTierUtilization node, the measured application tier utilization as an evidence for AppTierUtilization node, and so on for all other parameters. The evidences applied to these top-level nodes are propagated through the Bayesian network. It executes the probabilistic inference, taking into account the coordinated states of the entire multi-tier system. Based on the evidence applied to the top-level network nodes, WebTierUtilization node state is inferred as AT with 100% probability. Similarly, the state of other nodes and its probability will be determined by the probabilistic inference. Because of the probabilistic dependencies defined by the CPTs of the WebTierState and DBTierState nodes, both these nodes now have the Overloaded state with 100% probability. Similarly, the AppTierState node has two states, i.e., Underloaded and Normal each with 50% probability. This leads to the state of Shopping for the node SessionsToAccept. The type of the incoming session is applied as an evidence to the node SessionType, which gives the node the state of Shopping. The inference process is finished as the state of the AdmitSession node is inferred as Admit with 100% probability. This results in the incoming session of shopping type to be accepted. If the incoming session is of any other type, it would be rejected. We note that in this case the probability of admitting the session is 100%. But this may not be the case for other values of the evidence. As long as the probability of the SessionsToAccept node with the state Admit is above a threshold value (say 60%), the incoming session of any type will be accepted.

IV. PERFORMANCE EVALUATION

A. Experimental Setup

To evaluate the session-based admission control strategies, we built a simulation model for a three-tier e-commerce site. Without being affected by the methods of implementations, the simulator can effectively evaluate the performance of the admission control strategies by itself. It consists of a customer

TABLE III
EVIDENCES APPLIED TO THE BAYESIAN NETWORK.

Multi-tier System Parameter	Value
Utilization of the web tier	83%
Processing time of the web tier	85 ms
Utilization of the application tier	48%
Processing time of the application tier	27 ms
Utilization of the database tier	95%
Processing time of the database tier	89 ms
Incoming session type	Shopping

generator, a session generator, a request generator, a web server, an application server, and a database server.

As our previous work on session-based service differentiation on e-commerce servers [18], the session-based workload is generated following the guidelines provided by the TPC-W benchmark specification. A session belongs to one of the three distinct traffic mixes, Browsing, Shopping and Ordering. The session time for the session and think time between the interactions are generated by an exponential distribution with a given mean [2]. The utilization and the processing time of each interaction is derived from the WIRT (Web Interaction Response Time) based on the observation that different mixes pose varying load on the tiers. As others in [11], we assume that Browsing mix is database tier intensive, Shopping mix is application tier intensive and Ordering mix is web tier intensive. The resource demand of a request on the web tier, application tier and database tier is set as 20%, 30% and 50% of the overall resource demand corresponding to its WIRT. The Bayesian network used in the simulation was modeled with the Netica software [3].

B. Why Not Ordering Sessions First

One may argue that under heavy load conditions, intuitively only ordering sessions be accepted since they are more likely than other sessions to result in economic benefits for the e-commerce site. We executed experiments to support our counter-intuitive argument that accepting only ordering sessions under a heavy load condition is not the right way to utilize the resources in a multi-tier website. We applied a workload mix containing just Ordering sessions to the website with the MBAC approach. We compare the throughput results with the experiment that applied a workload mix consisting of an equal number of sessions of all three mixes.

Figures 5(a)(b)(d)(e) show the experimental results of the number of accepted sessions, completed sessions, rejected sessions, and aborted sessions, respectively. Note that we represent the workload as the number of arrival sessions per second instead of a utilization percent. This is due the fact that there is no single utilization percent in a multi-tier architecture and the bottleneck tier shifts from time to time. The results show that accepting different traffic mixes can significantly increase the throughput, both as the accepted sessions and as the completed sessions. Figures 5(c)(f) depict the improvement in percentages. The rationale is that just accepting one type of sessions will overload one of the tiers sooner, leading to

Fig. 4. Predicted utilizations of a three-tier website.

Fig. 5. Throughput comparison with varying traffic mix.

the system imbalance. On the other hand, accepting a mixture of different sessions will place more balanced utilizations across all three tiers. Thus the effective session throughput is increased significantly.

C. Impact of CoSAC on Effective Session Throughput

We now evaluate the impact of CoSAC with a Bayesian network on the session throughput and compare it with the results due to the use of Blackbox and MBAC approaches. The workload consists of an equal number of Browsing, Shopping and Ordering sessions. We conducted experiments at different session arrival rates between 10 to 100 sessions/sec. Figures 6(a)(b)(d)(f) show the number of accepted, rejected, completed, and aborted sessions, respectively. Results demonstrate that CoSAC is able to accept and complete significantly more sessions than the MBAC and Blackbox approaches. When the overall session arrival rate is greater than 80 sessions/sec, the saturation point of the website is

reached. Using the results due to the Blackbox as the baseline, Figure 6(c) shows that CoSAC is able to accept as many as 45% more sessions and MBAC is able to accept about 20% more sessions. Figure 6(f) shows that CoSAC is able to complete as many as 50% more sessions and MBAC is able to complete as many as 22% more sessions. This demonstrate the significance of multi-tier admission control approaches. Particularly, the admission coordination between multiple tiers via the Bayesian network representation leads to significant session throughput improvement. We performed a wide range of sensitivity analysis. All results demonstrate the superior performance of CoSAC, which uses a Bayesian network to coordinate the session admission in a multi-tier website. It outperforms other two approaches significantly with respect to the number of completed sessions (the effective session throughput). Due to the space limitation, we only presented those representative results. For more details, please refer to the comprehensive technical report [10].

Fig. 6. Impact of admission control strategies on session throughput (with 1:1:1 workload mix).

V. CONCLUSION

In this paper, we proposed and evaluated two new session-based admission control approaches for multi-tier Internet applications. Both recognize the fact that the bottleneck in a multi-tier website shifts among tiers as client access patterns change dynamically. The existing session-based admission control approach (SBAC) is based on the bottleneck resource, and thus not effective as different client sessions impose different resource consumptions at the different tiers. The new MBAC approach is based on the dynamic measurement of utilizations at the different tiers. It makes admission decision pro-actively based on the session mix characteristics. The new CoSAC approach novelly uses a machine learning technique, a Bayesian network, to coordinate the admission decision based on the utilizations of the multiple tiers. The research work is significant. Extensive experimental results using TPC-W benchmark workloads demonstrate the superior performance of the new approaches. Particularly, CoSAC can improve the effective session throughput by about 50% compared to the SBAC-based approach for multi-tier Internet applications.

Acknowledgement

This research was supported in part by U.S. National Science Foundation under Grant No. CNS-0720524.

REFERENCES

- [1] J. Chen, G. Soundararajan, and C. Amza. Autonomic provisioning of backend databases in dynamic content Web servers. In *Proc. IEEE ICAC*, 2006.
- [2] L. Cherkasova and P. Phaal. Session-based admission control: A mechanism for peak load management of commercial web sites. *IEEE Trans. on Computers*, 51(6):669–685, 2002.
- [3] N. S. Corporation. Netica-j: Java netica api. <http://www.norsys.com/netica-j.html>.
- [4] Y. Diao, J. L. Hellerstein, S. Parekh, H. Shaihk, and M. Surendra. Controlling quality of service in multi-tier Web applications. In *Proc. IEEE ICDSC*, 2006.
- [5] S. Elnikety, E. Nahum, J. Tracey, and W. Zwaenepoel. A method for transparent admission control and request scheduling in e-commerce web sites. In *Proc. of ACM WWW*, 2004.
- [6] J. Guitart, D. Carrera, V. Beltran, J. Torres, and E. Ayguade. Session-based adaptive overload control for secure dynamic web applications. In *Proc. IEEE ICPP*, 2005.
- [7] T. Horvath, T. Abdelzaher, K. Skadron, and X. Liu. Dynamic voltage scaling in multitier web servers with end-to-end delay control. *IEEE Trans. on Computers*, 56(4):444–458, 2007.
- [8] A. Kamra, V. Misra, and E. M. Nahum. Yaksha: a self-tuning controller for managing the performance of 3-tiered web sites. In *Proc. IEEE IWQoS*, 2004.
- [9] P. Lama and X. Zhou. Efficient server provisioning with delay guarantee on multi-tier clusters. In *Proc. IEEE IWQoS*, 2009.
- [10] S. Muppala and X. Zhou. CoSAC: coordinated session admission control with statistical learning for multi-tier Internet applications. In *Technical Report, Computer Science Department, UCCS*, 2009.
- [11] J. Rao and C. Xu. CoSA: a coordinated statistical learning approach to measuring the capacity of multi-tier Websites. In *Proc. IEEE IPDPS*, 2008.
- [12] B. Urgaonkar, G. Pacific, P. Shenoy, M. Spreitzer, and A. Tantawi. An analytical model for multi-tier Internet services and its applications. In *Proc. ACM SIGMETRICS*, 2005.
- [13] B. Urgaonkar, P. Shenoy, A. Chandra, P. Goyal, and T. Wood. Agile dynamic provisioning of multi-tier Internet applications. *ACM Trans. on Autonomous and Adaptive Systems*, 3(1):1–39, 2008.
- [14] D. Villela, P. Pradhan, and D. Rubenstein. Provisioning servers in the application tier for e-commerce systems. *ACM Trans. on Internet Technology*, 7(1):1–23, 2007.
- [15] J. Wildstorm, P. Stone, E. Witchel, and M. Dahlin. Machine learning for on-line hardware reconfiguration. *International Joint Conferences on Artificial Intelligence*, 14(2):108–136, 2007.
- [16] Q. Zhang, L. Cherkasova, and E. Smirni. A regression-based analytic model for dynamic resource provisioning of multi-tier Internet applications. In *Proc. IEEE ICAC*, 2007.
- [17] J. Zhou and T. Yang. Selective early request termination for busy internet services. In *Proc. ACM WWW*, 2006.
- [18] X. Zhou, J. Wei, and C.-Z. Xu. Resource allocation for session-based two-dimensional service differentiation on e-commerce servers. *IEEE Trans. on Parallel and Distributed Systems*, 17(8):838–850, 2006.