

Automated Cluster-Based Web Service Performance Tuning

I-Hsin Chung, and Jeffrey K. Hollingsworth
{ihchung, hollings}@cs.umd.edu

Department of Computer Science
University of Maryland
College Park, MD 20742

Abstract – Active Harmony provides a way to automate performance tuning. In this paper, we apply the Active Harmony system to improve the performance of a cluster-based web service system. The performance improvement cannot easily be achieved by tuning individual components for such a system. The experimental results show that there is no single configuration for the system that performs well for all kinds of workloads. By tuning the parameters, Active Harmony helps the system adapt to different workloads and improve the performance up to 16%. For scalability, we demonstrate how to reduce the time when tuning a large system with many tunable parameters. Finally an algorithm is proposed to automatically adjust the structure of cluster-based web systems, and the system throughput is improved up to 70% using this technique.

I. INTRODUCTION

Online e-commerce sites are one of the main applications on the Internet today. They are used as a standard mechanism for online information distribution and exchange. In order to provide such service, e-commerce sites require large online web systems. The systems must be able to accommodate widely varying service demands. They should also be adaptive when the number or nature of requests changes.

Clusters of commodity workstations interconnected by a high-speed network are frequently used to meet these challenges. The infrastructure can tolerate partial failures and allows scaling up by adding more components. They are also representative of other types of coupled distributed systems.

When these systems are designed and built, the developers tend to set the default configuration of the system (e.g., number of processes forked, memory size allocated) conservatively (i.e., appropriate values but not well tuned). Therefore, the customer environment may not be fully utilized and thus the performance for such a system may be improved if its configuration is “tuned” appropriately.

While other clustered-based web service performance tuning projects require experts to analyze the internals of the components and improve the performance based on the models built, the Active Harmony system is designed to provide a general solution that can help systems become adaptive to their execution environment as well as to changes in workload. By improving the performance iteratively, the Active Harmony system changes performance optimization

from post-mortem to real-time steering. And the most important of all, it is not necessary for the Active Harmony user to have detailed insight knowledge of the system to be tuned.

This paper differs from our previous work [9, 11, 20] in that we propose *parameter replication* and *parameter partitioning* to speed up the tuning process. We also present and evaluate a technique to allow Active Harmony to reconfigure the roles of specific nodes during execution. We then apply Active Harmony to a coupled application. An e-commerce system contains multiple components (proxy server, HTTP server, application server, and database). Such a large-scale system cannot be tuned for each individual component. In this paper we show that Active Harmony is not only useful to improve the performance, but it is necessary to have such a tuning mechanism since there is no single best configuration for all kinds of workloads.

II. SYSTEM

A cluster-based web service system consists of a collection of machines. The machines are separated into sets. Each set (or tier) of machines is focused on serving different parts of a request. The incoming requests are handled in a pipeline fashion by different tiers.

In many web services today, there are (conceptually, at least) three tiers: presentation, middleware, and database. The presentation tier is the web server that provides the interface to the client. The middleware tier is what sits between the web server and the database. It receives requests for data from the web server, manipulates the data and queries the database. Then it generates results using existing data together with answers from database. Those results are presented to the client through the presentation tier. The third tier is the database, which holds the information accessible via the Web. It is the backend that provides reliable data storage and transaction semantics.

In this project, we try to improve the overall system performance by automatic tuning across all tiers using the Active Harmony system. The performance metric we are focusing on is the TPC-W benchmark. It is a transactional web benchmark designed to emulate operations of an e-commerce site.

A. Active Harmony

To provide automatic performance tuning, we developed the Active Harmony system [9, 11, 20]. Active Harmony is an infrastructure that allows applications to become tunable by applying very minimal changes to the application and library source code. This adaptability provides applications with a way to improve performance during a single execution based on the observed performance. The types of things that can be tuned at runtime range from parameters such as the size of a read-ahead parameter to what algorithm is being used (e.g., heap sort vs. quick-sort).

Figure 1 shows the Active Harmony automated runtime tuning system. The Library Specification Layer provides a uniform API to library users by integrating different libraries with the same or similar functionality.

The Adaptation Controller is the main part of the Harmony server. The Adaptability component manages the values of the different tunable parameters provided by the applications and changes them for better performance.

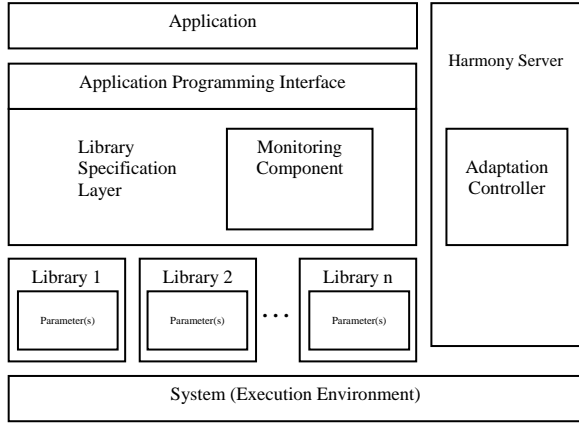


Figure 1: Active Harmony automated tuning system

The kernel of the adaptation controller is a tuning algorithm. The algorithm is based on the simplex method for finding a function's minimum value [14]. In the Active Harmony system, we treat each tunable parameter as a variable in an independent dimension. The algorithm makes use of a simplex, which is a geometrical figure defined by $k+1$ connected points in a k -dimensions space. In 2-dimensions, the simplex is a triangle, and for 3-d space the simplex is a non-degenerated tetrahedron.

The Nelder-Mead simplex method approximates the extreme of a function by considering the worst point of the simplex and forming its symmetrical image through the center of the opposite (hyper) face. At each step a better point replaces the worst points and thus moves the simplex towards the extreme, in our case towards the minimum.

The algorithm described above assumes a well-defined function and works in a continuous space. However, neither of these assumptions holds in our situation. Thus we have adapted the algorithm by simply using the resulting values

from the nearest integer point in the space to approximate the performance at the selected point in the continuous space.

B. TPC-W Benchmark

The major workload we use when tuning the cluster-based web service is the TPC-W benchmark. The TPC-W is a transactional web benchmark designed to mimic operations of an e-commerce site. The workload explores a breadth of system components together with the execution environment. Like all other TPC benchmarks, the TPC-W benchmark specification is a written document which defines how to setup, execute, and document a TPC-W benchmark run.

TABLE 1: TPC-W BENCHMARK WORKLOADS

Web Interaction	Browsing (WIPSB)	Shopping (WIPS)	Ordering (WIPSO)
Browse	95 %	80 %	50 %
Home	29.00 %	16.00 %	9.12 %
New Products	11.00 %	5.00 %	0.46 %
Best Sellers	11.00 %	5.00 %	0.46 %
Product Detail	21.00 %	17.00 %	12.35 %
Search Request	12.00 %	20.00 %	14.53 %
Search Results	11.00 %	17.00 %	13.08 %
Order	5 %	20 %	50 %
Shopping Cart	2.00 %	11.60 %	13.53 %
Customer Registration	0.82 %	3.00 %	12.86 %
Buy Request	0.75 %	2.60 %	12.73 %
Buy Confirm	0.69 %	1.20 %	10.18 %
Order Inquiry	0.30 %	0.75 %	0.25 %
Order Display	0.25 %	0.66 %	0.22 %
Admin Request	0.10 %	0.10 %	0.12 %
Admin Confirm	0.09 %	0.09 %	0.11 %

The two primary performance metrics of the TPC-W benchmark are the number of Web Interaction Per Second (WIPS), and a price performance metric defined as Dollars/WIPS. However, some shopping applications attract users primarily interested in browsing, while others attract those planning to purchase. Two secondary metrics are defined to provide insight as to how a particular system will perform under these conditions. WIPSB is used to refer to the average number of Web Interaction Per Second completed during the Browsing Interval. WIPSO is used to refer to the average number of Web Interaction Per Second completed during the Ordering Interval.

The TPC-W workload is made up of a set of web interactions. Different workloads assign different relative weights to each of the web interactions based on the scenario. In general, these web interactions can be classified as either "Browse" or "Order" depending on whether they involve browsing and searching on the site or whether they play an explicit role in the ordering process. The details for each workload breakdown are shown in the Table 1.

C. Environment

The summary of the environment used for our experiment is shown in Table 2. The 10 machines used include the ones running emulated browsers and the servers for proxy, HTTP, application and database services. Each machine is equipped with dual processors, 1 Gbyte memory and runs Linux as the operating system. For each tier, we select Squid as the proxy server, Tomcat as the HTTP & application server and MySQL as the database server. All computer software components are open-source which allows us to look at source code to understand system performance. The TPC-W benchmark version we chose simulates a store that sells approximately 10,000 items.

TABLE 2: EXPERIMENT ENVIRONMENT

Hardware	
Processor	Dual AMD Athlon 1.67 GHz
Memory	1Gbyte
Network	100Mbps Ethernet
No. of machines	10
Software	
Operating System	Linux 2.4.18smp
TPC-W benchmark	Modified from the PHARM [6]
Proxy Server	Squid 2.5 [3]
HTTP & Application Server	Tomcat 4.0.4 [1]
Database Server	MySQL 3.23.51 [2]

III. TUNING

Our goal is to improve the overall system performance using Active Harmony. We first show that there is no single configuration suitable for all the workloads. Active Harmony makes the system perform better by using different configurations when facing different workloads. Then we investigate Active Harmony’s scalability as the number of machines grows. One way to solve this problem is to partition the parameters into sets. We show how to use an independent Active Harmony tuning server for each set to speed up the tuning process. Another method is to tune a representative set of parameters and use duplicated values on the rest of nodes. In Section four, we also show how to adjust the number of nodes in each tier dynamically to reduce hot spots.

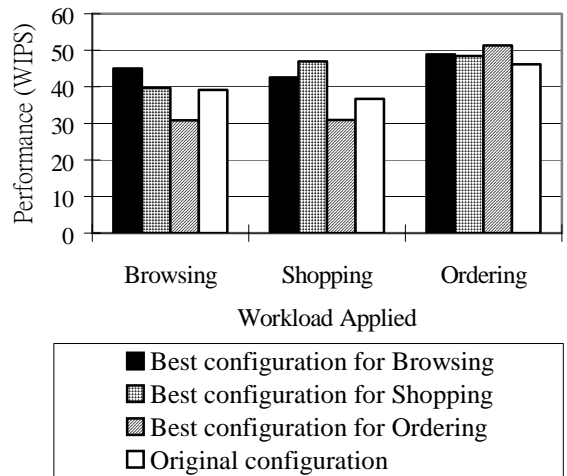
A. Impact of Varying Workload

In this experiment we show that the Active Harmony server can tune the system to adjust each tier’s server to provide good performance. We use four machines in this experiment: one machine for the emulated browsers, one for the proxy server, one for the HTTP & application server, and one for the database server.

In the experiment, we examine the tuning processes for two different workloads: browsing and ordering. Both tuning processes are started using the default configuration. We then let the system warm up for 100 seconds and measure the performance (WIPS) for 1000 seconds followed by 100 seconds for cooling down. We define such a cycle as one

“iteration”¹. The Active Harmony server will adjust the configuration (parameters values) between two iterations.

Figure 2 shows that for different workloads, the system should apply different configurations. Each different bar represents the best configurations we determined after 200 tuning iterations for each of the workloads. We then apply those best configurations to the other two workloads for comparison. The results show that when using a configuration that is tuned for another workload, the system does not perform as well as using a configuration that is tuned for the current workload. The results show that there is no universal configuration good for all kinds of workloads. The table in Figure 2 shows the improvements for those best-tuned configurations compared to the default configuration. The improvements range from 5% to 16%.



	Best configuration after 200 iterations		
	Browsing	Shopping	Ordering
Improvement compared to the default configuration	15%	16%	5%

Figure 2: Applying best configuration after 200 iterations to different workloads

Table 3 shows the details of all Harmony tunable parameters before, and after tuning for each of the workloads. The results show for the proxy server, it first increases the main memory size for the cache to improve the performance. For the shopping and ordering workloads, the proxy server tries to cache larger objects in the memory compared to the browsing workload. For the HTTP server (which is part of the application server), the tuning results show that it spawns more threads to handle the requests during the ordering workload. We believe the main reason is that most of the requests in the ordering workload require high latency operations in the database server (i.e., performing update

¹ The 1,200 second-iteration is TPC-W benchmark compliance (i.e., specified in the TPC-W documentation). The iteration timescale can be as short as 30 seconds according to our experiment experience.

transactions on the database). Thus the average response time is longer compared to other workloads. As long as it is not over the system capacity, the HTTP server should use more threads (minProcessors/maxProcessors) and buffer space (bufferSize) to handle the incoming requests. The waiting queue capacity should also increase accordingly (acceptCount) as the results show. The same situation happens in the worker part (AJP connector) of the application server. For the database server, the tuning results show it increases the cache and buffer size when the utilization for the database is high (i.e., shopping and ordering workloads). However, it shows that reducing the join buffer size does not impact performance.

From the results we can see that some parameters significantly affect the overall system performance such as the number of threads or the buffer size. However, there are some parameters that we thought to be performance related but they turn out not to be important. For example, the thresholds (cache_swap_low, cache_swap_high) which control whether the proxy server should swap out objects do not impact the overall system performance. Since it is automated, the Active Harmony tuning process is also helpful for system administrators and developers to identify those parameters that actually affect system performance. We plan to further address this issue by prioritizing the importance of parameters in our future work.

TABLE 3: TUNING RESULTS FOR DIFFERENT WORKLOADS

Tunable parameters	Default config.	Best configuration after 200 iterations		
		Browsing	Shopping	Ordering
Proxy Server				
cache_mem	8	13	17	21
cache_swap_low	90	91	86	91
cache_swap_high	95	96	96	96
maximum_object_size	4,096	4,096	4,096	5,888
minimum_object_size	0	0	50	306
maximum_object_size_in_memory	8	6	256	2,560
store_objects_per_bucket	20	15	25	105
HTTP & App. Server				
minProcessors	5	1	16	102
maxProcessors	20	11	16	131
acceptCount	10	6	21	136
bufferSize	2,048	2,049	3,585	6,657
AJPminProcessors	5	6	26	136
AJPmaxProcessors	20	86	296	161
AJPacceptCount	10	76	306	671
Database Server				
binlog_cache_size	32,768	63,488	153,600	284,672
Delayed_insert_limit	100	200	400	700
max_connections	100	201	451	701
delayed_queue_size	1000	2,600	9,100	7,100
Join_buffer_size	8,388,600	407,552	407,552	407,552
Net_buffer_length	16,384	31,744	38,912	34,816
table_cache	64	873	905	761
thread_con	10	81	91	76
thread_stack	65,535	102,400	1,018,880	773,120

B. Cluster Tuning

When the number of servers increases, the number of tunable parameters also increases. This makes the tuning

process lengthy and the tuning results may not be useful since the environment could change during the tuning process.

In the original Active Harmony system, to tune n parameters at once requires exploring $n+1$ configurations before improvements to the system will take effect. If there are numerous servers in the cluster and each server contains tens of parameters, the tuning process will be fairly long. In order to reduce the initial exploration period, we partition the components inside the cluster into groups and use separate Active Harmony tuning servers for each groups. There are several ways to group servers.

When all the machines in the same tier are homogeneous, we try to partition all the servers into tuning groups using two methods. The first one is parameter duplication: we only tune one server for each tier, and the values for those parameters are duplicated to other servers in the same tier. This tuning mechanism is based on the assumptions that (a) servers in the same tier will have the same or similar behavior for the same configuration; (b) the workload is evenly distributed among all the servers in the same tier.

The second way to group nodes, parameter partitioning, is based on a static work line. Each work line group consists of at least one server from each tier. A request to the web cluster system is only handled by exactly one work line group. In other words, any server in work line group A will not generate (serve) requests to (from) a server in work line group B. We use a different Active Harmony tuning server to tune the parameters for each work line. The assumption for this tuning mechanism is that (a) all the work lines are running in parallel; and (b) there is no interaction between any two of the work lines.

Both of these approaches to grouping nodes require some domain knowledge about the role of each node. However, grouping of nodes could easily be exported to Active Harmony as part of the tuning API.

To compare these two approaches, we tuned the system using three different tuning methods: default, parameter duplication and parameter partitioning.

TABLE 4: PERFORMANCE FOR DIFFERENT METHODS FOR CLUSTER TUNING

Tuning method	WIPS ²	Average (Std. Dev.) ³	Performance improvement	Iterations
None (No Tuning)	110.4	110.4 (2.1)	-	-
Default method	130.6	112.1 (30.0)	18.3%	159
Parameter duplication	133.7	116.6 (29.5)	21.2%	33
Parameter partitioning	131.3	121.8 (9.7)	19.0%	107

Table 4 shows the tuning results. The results for all three methods are very close. The default method takes the longest time since there are many parameters and only one

² Performance for the best configuration after 200 iterations

³ For the second 100 iterations

performance result per iteration. The parameter duplication method provides both a larger performance improvement and faster convergence to the tuned configuration. It speeds up the tuning process since the tunable parameters are distributed to multiple tuning servers and there are fewer parameters for each tuning server to tune. The time (iterations) spent for the grouping by parameter partitioning method is about 2/3 of the default method.

Based on the time for the tuning process, parameter duplication tuning seems to be the best. It takes a much shorter time for tuning. However, if stable performance during the tuning process is critical, parameter partitioning by work lines is a reasonable choice.

In the future, we plan to investigate hybrid tuning using the parameter duplication method first, and then using a separate tuning server for each group for fine-granularity tuning.

IV. AUTOMATIC CLUSTER RECONFIGURATION

One of the advantages for a cluster-based web service is the ability to reconfigure hardware easily. By dynamically changing the roles of servers for different workloads, it is possible to make the best of available resources.

The parameter tuning part of the Active Harmony system helps to tune the cluster-based web service at a fine time granularity. However, when the load is not balanced among tiers in the web service system, changing the parameters for all the servers will not provide much help to solve the problem. Instead, it is necessary to adjust the infrastructure by changing the number of servers in each tier dynamically to reduce the load imbalance.

TABLE 5: VARIABLE DESCRIPTION

Variable	Description
R_{ij}	Utilization of resource j on node i
LT_{ij}	Low threshold for resource j on node i
HT_{ij}	High threshold for resource j on node i
M_{pq}	Cost to move a job for node p to node q
A_i	Average process time on node i
F	Configuration cost in terms of time
L	List of nodes
N_i	Number of jobs on node i
$Head(L)$	First node in the List L
$Tier(i)$	The tier that node i belongs to
$M(t)$	Number of nodes in tier t

1. For all node i , resource j do
 If $R_{ij} > HT_{ij}$ then add i to the list L_1
 //find out what nodes are highly or over loaded
2. For all node i do
 If $R_{ij} < LT_{ij}$ for all j then add i to the list L_2
 //find out what nodes are lightly loaded
3. Sort L_1 based on the “degree of urgency”⁴

⁴ The degree of urgency for each node depends on the characteristics of the application. It may vary from case to case. For example, over loading the CPU may cause bigger problem than utilizing all the network bandwidth for

- //decide the priority for the nodes to be relieved
4. Let $i = Head(L_1)$, find the node k in L_2 such that satisfies (a)(b)(c)
 //find out the appropriate node to be reconfigured
 (a) $Tier(i) \neq Tier(k)$
 (b) $M(Tier(k)) > 1$
 (c) $F + N_k \times M_{km} - N_k \times A_k$ is minimal, where $k \neq m$
 and $Tier(k) = Tier(m)$
 5. Reconfigure k such that $Tier(i) = Tier(k)$

Figure 3: Reconfiguration algorithm for external tuning

The Active Harmony system applies a simple mechanism to achieve load balance among tiers. While the tuning is in progress, the Active Harmony system monitors the resource utilization for all nodes of all tiers. The resources that are monitored include CPU load, memory usage, network bandwidth used and disk I/O activity (Currently the system information is obtained using Linux SAR utility tool). Periodically, Active Harmony detects whether (1) there is a resource on node A that is over utilized⁵, (2) all the resources on node B are under utilized and node B is suitable reconfiguration. If both situation (1) and (2) exist, Active Harmony tries to reconfigure node B to run the same server process as node A.

Unlike parameter tuning, which is done for each iteration, the reconfiguration algorithm is run at a lower frequency (e.g., every 50 iterations) since it is designed to react to longer term trends, and incurs a greater overhead to make changes. Table 5 shows the definition for variables in the algorithm and Figure 3 shows the concept of the reconfiguration algorithm.

Step 1 finds out what nodes are over loaded. It checks the resource utilization against the predefined high threshold. Step 2 tries to find nodes that are lightly loaded. If all the resources on the node are idling most of the time (i.e., utilization is smaller than the lower threshold), the node is considered under utilized. Step 3 finds out what is the most “urgent” node that should be relieved first. Step 4 checks in order to ensure correct operation, that there is at least one node left in each tier, and decides if the reconfiguration should be done immediately (by moving existing requests to the neighbor nodes in the same tier) or if it should wait until all existing requests finish. Finally Step 5 does the reconfiguration.

$$F + N_k \times M_{km} - N_k \times A_k \quad (1)$$


When the result of equation (1) for the selected node k in Step 4(c) is non-negative, the Active Harmony system will not reconfigure node k until all the jobs on it are finished. This is because it will be more cost-effective to wait than to reconfigure node k immediately. On the other hand, when the result of the equation is negative, the Active Harmony system

some applications. Therefore, nodes with over-loaded CPU will have higher priority than nodes whose network bandwidth is highly utilized.


⁵ Static thresholds (e.g., CPU idle time is less or equal than 5%) are used in the current implementation.

will reconfigure node k immediately. This is because the cost for immediate reconfiguration will be less than waiting for the system to be idle to reconfigure.

Active Harmony can automatically perform node reconfiguration without taking the system down. While one node is being reconfigured from one tier to another, all the remaining nodes in the system are still serving requests normally.



(a) One node moved from the proxy server tier to the application server tier
(Workload changes from browsing to ordering)



(b) One node moved from the application server tier to the proxy server tier
(Browsing workload)

Figure 4: Reconfiguration experiment results

Figure 4 shows the experimental results when applying the reconfiguration algorithm. The initial configuration for Figure 4(a) has four nodes serving the proxy tier and another two nodes for the application tier; all six nodes are homogeneous. The experiment starts with a browsing workload and changes to an ordering workload after the 90th iteration (The performance gains between 90th and 100th iterations are due to different workloads). We forced the Active Harmony system do the dynamic adjustment checking exactly once right after the 100th iteration of the tuning process. Figure 4(a) shows the performance improvement

when Active Harmony decides to move a node from the proxy server tier to the application server tier based on the algorithm. This is expected since when the system has a workload dominated by ordering, it requires more application servers to handle the dynamic data from the database. On the other hand, most browsing workloads require static data that can be served from the proxy servers. Before the adjustment, the application servers are highly loaded (CPU utilization is always close to 100%) and some proxy servers are idling most of the time (CPU utilization is close to 0% and there are very few network or disk I/O requests). After the adjustment, the average utilization of the application servers is lowered while the average loading for the proxy servers increases a little. The bottleneck of the whole system is relieved and the system performance is improved about 62%.

Figure 4(b) shows the performance improvement when given a different configuration at the beginning. There are six nodes, two of them serving as the proxy servers and four serving as application nodes. However, the proxy servers are highly utilized under the browsing workload. After the dynamic adjustment checking after the 100th iteration, it moved a node from the application server tier to the proxy server tier for the adjustment automatically. The CPU and disk I/O are highly loaded on the proxy servers before the adjustment and some application servers are idling most of the time. After the adjustment, the average load on all proxy servers is lowered, the average utilization on the remaining application servers is increased and the system performance is improved for about 70%.

V. DISCUSSION

To tune existing software such as the Squid proxy server, we needed to make some minimal modifications to add calls to the Active Harmony API. However, some variables are only referenced once after the program starts execution (i.e., those variables read from the configuration script file). Rather than make more extensive changes to the program, the Active Harmony system restarts the server for each of the tuning iterations automatically. Our experiments take all costs of parameter changes (including servers need to be restarted and their warm up time) into consideration.

Another issue is the hard coded (compile time) limits in the applications. In order to make the system tunable, some limits had to be increased. Again, a more significant coding effort could have been used to convert these hard-coded limits into ones that could be changed at runtime. For example, to increase the number of files opened simultaneously, the value in the `/proc/sys/fs/file-max` on Linux needed to be increased. Otherwise the number of files opened simultaneously would be limited. In this case, recompilation of the linux kernel would be necessary. Besides the kernel, the linux operating system also imposes similar constraints in the `/etc/security/limits.conf` and `/etc/sysctl.conf`.

Active Harmony helps the cluster-based web service adapt itself when facing different workloads. It shows the ability to tune a large-scale system automatically. The tuning includes the parameter adjustment inside each machine and

explicit configuration changes for load balancing. This performance improvement is difficult to achieve by tuning each single machine independently since it is extremely difficult to decide the contribution of each individual machine to the performance of the whole system. Another advantage is that the user does not need to have detailed insight knowledge about each component. He or she can simply apply the Active Harmony system to all the parameters that may be performance related.

VI. RELATED WORK

There are several projects that are trying to develop techniques to allow applications to be responsive to their available resources or that allow them to be tuned at runtime. The Falcon project [8] focuses on computational steering. It provides a way for users to alter the behavior of an application under execution. The execution results are also changed based on the steering mechanism. The Active Harmony project also allows user to alter the configuration during execution but it is focusing on performance tuning rather than the experiment result.

The Autopilot project [16, 17] allows applications to be adapted in an automated way. It uses sensors to extract quantitative and qualitative performance data from executing applications, and provides the requisite data for decision-making. The kernel of the decision process for Autopilot is fuzzy logic. Their actuators execute the decision by changing parameter values of applications or resource management policies of the underlying system. The Active Harmony project differs from the Autopilot project in that it tries to coordinate the use of resources by multiple libraries and applications rather than focusing on a single application.

The AppLes project [5] and the Odyssey project [15] focus on resource awareness at the application level. In those systems, applications are informed of resource changes and provided with a list of available resource sets. Then, each application allocates the resources based upon a customized scheduling to maximize its own performance. Active Harmony encourages programmers to expose their needs in terms of options and their characteristics rather than as selecting from specific resources alternatives described by the system.

The ATLAS [21] project has developed automatically tuned linear algebra libraries. They develop a methodology for the automatic generation of highly efficient basic linear algebra routines for a given microprocessor. By using a code generator that probes and searches the system for an optimal set of parameters, it can produce highly optimized matrix multiply for a wide range of architectures. The difference between our work and ATLAS is that our work focuses on general applications that use program libraries rather than that of a specific library.

The Nimrod/O project [4] tries to reduce the search space for engineering design. It applies multiple tuning algorithms including Simplex, P-BFGS, Divide and Conquer, and Simulated Annealing. The design for the aerofoil may need to search for the global optima instead the local optima. The

Active Harmony project focuses on the performance issue. Therefore, operating points on local optima are still acceptable in most of the cases since they are also good enough from the performance point of view.

Another TPC-W benchmark implementation available from an academic institute is from the DynaServer project [19]. The project studies the design of scalable, high-performance and highly available e-business servers.

Others have discussed cluster-based web services with different performance metrics. Joel L. Wolf's work [22] proposed a scheme, which attempts to optimally balance the load on the servers of a clustered Web farm. They try to solve the performance problem by achieving minimal average response time for customer requests, and thus ultimately achieve maximal customer throughput.

ADAPTLLOAD [18] developed by Riska, A., et al. models clustered web server as a front-end dispatcher and back-end nodes. They use an online algorithm to decide the share of the total workload for each node to achieve load balance. They treat back-end nodes as static while Active Harmony tries to configure the clustered system properly to achieve better performance.

Chen, et al. [7] use a reconfiguration mechanism to improve the throughput of a clustered system. Their focus is to avoid letting a small number of running jobs with unexpectedly large memory allocation block the execution of the majority jobs in the cluster. Active Harmony focuses on a general mechanism to improve overall system performance by several means.

Kalogeraki, et al. [10] migrate objects or jobs from hotspots in the cluster to improve the performance. Their goal is to achieve load balance while Active Harmony focuses on performance improvement.

Gage [13] focuses on load distribution to provide the performance guarantee for cluster-based Internet services. This involves support from network level while the Active Harmony only tries to tune the system to achieve better performance.

Levy, et al. [12] use a queuing model to analyze a cluster-based web service system. Based on the model built, they implement a prototype for a performance management system that is transparent to the system to be tuned.

The major difference between Active Harmony and these works above is that Active Harmony provides a general solution that does not require the user to have domain specific knowledge. The user does not need to analyze the details of the system components or build models.

VII. CONCLUSION

The main contribution for this paper is that we apply Active Harmony to a coupled system of independent applications. We applied Active Harmony to a real-world large-scale system and evaluated the result using a practical benchmark. The tuning includes the parameters adjustment inside each machine and the explicit configuration change for load balancing. All this is done without the user needing to have domain specific information.

The performance improvement is difficult to achieve when tuning individual components of the system separately. Since no single universal configuration is good for all kinds of workloads, the cluster based web service system needs a tuning mechanism like the Active Harmony. Active Harmony adjusts the tunable parameters based on the observed performance results to improve the overall system performance. The experiment results show that the Active Harmony system improves the system performance from 5% to 16% depending on the workload.

Scalability becomes a critical issue when tuning large-scale systems with numerous parameters. We investigated two approaches for tuning – parameter replication and parameter partitioning. This is helpful to speed up the tuning process so the tuning results will not be out of date. Parameter duplication helps to speedup the tuning process while parameter partitioning makes the tuning process smoother with stable performance.

Dynamically adjusting the components of the cluster, the performance is improved by better load balancing. In our experiments, the system throughput is improved up to 70%. All the results demonstrate that Active Harmony can bring significant performance improvement to the cluster-based web service system and permit new ways to adapt applications to dynamic environments.

ACKNOWLEDGEMENT

This work was supported in part by NSF award EIA-0080206 and DOE Grant DE-FG02-01ER25510.

REFERENCES

1. *The Apache Jakarta Project* <http://jakarta.apache.org/>.
2. *MySQL Database Server*, MySQL AB <http://www.mysql.com>.
3. *Squid Web Proxy Cache* <http://www.squid-cache.org/>.
4. Abramson, D., et al. *An Automatic Design Optimization Tool and its Application to Computational Fluid Dynamics*. in *SC*. 2001. Denver.
5. Berman, F. and R. Wolski. *Scheduling from the perspective of the application*. in *Proceedings of 5th IEEE International Symposium on High Performance Distributed Computing*. 1996. Syracuse, NY, USA 6-9 Aug. 1996.
6. Bezenek, T., et al., *Java TPC-W Implementation Distribution* <http://www.ece.wisc.edu/~pharm/tpcw.shtml>.
7. Chen, S., L. Xiao, and X. Zhang. *Adaptive and Virtual Reconfigurations for Effective Dynamic Job Scheduling in Cluster Systems*. in *22 nd International Conference on Distributed Computing Systems (ICDCS'02)*. 2002. Vienna, Austria.
8. Gu, W., et al. *Falcon: On-line Monitoring and Steering of Large-Scale Parallel Programs*. in *Frontiers '95*. 1995. McLean, VA: IEEE Press.
9. Hollingsworth, J.K. and P.J. Keleher. *Prediction and Adaptation in Active Harmony*. in *The 7th International Symposium on High Performance Distributed Computing*. 1998. Chicago.
10. Kalogeraki, V., P.M. Melliar-Smith, and L.E. Moser. *Dynamic Migration Algorithms for Distributed Object Systems*. in *The 21st International Conference on Distributed Computing Systems*. 2001. Mesa, AZ.
11. Keleher, P.J., J.K. Hollingsworth, and D. Perkovic. *Exposing Application Alternatives*. in *ICDCS*. 1999. Austin, TX.
12. Levy, R., et al. *Performance Management for Cluster Based Web Services*. in *The 8th IFIP/IEEE International Symposium on Integrated Network Management (IM2003)*. 2003. Colorado Springs, Colorado, USA.
13. Li, C., et al. *Performance Guarantee for Cluster-Based Internet Services*. in *The 23rd IEEE International Conference on Distributed Computing Systems (ICDCS 2003)*. 2003. Providence, Rhode Island.
14. Nelder, J.A. and R. Mead, *A Simplex Method for Function Minimization*. *Comput. J.*, 1965. **7**(4): p. 308--313.
15. Noble, B.D., et al. *Agile Application-Aware Adaptation for Mobility*. in *16th ACM Symposium on Operating Systems Principals*. 1997.
16. Ribler, R.L., H. Simitci, and D.A. Reed, *The Autopilot Performance-Directed Adaptive Control System*. *Future Generation Computer Systems*, special issue (Performance Data Mining), 2001. **18**(1): p. 175-187.
17. Ribler, R.L., et al. *Autopilot: Adaptive Control of Distributed Applications*. in *High Performance Distributed Computing*. 1998. Chicago, IL.
18. Riska, A., et al. *ADAPTLOAD: Effective Balancing in Clustered Web Servers Under Transient Load Conditions*. in *22 nd International Conference on Distributed Computing Systems (ICDCS'02)*. 2002.
19. Snavely, A., et al. *A Framework for Application Performance Modeling and Prediction*. in *Supercomputing 2002*. 2002. Baltimore, MD.
20. Tapus, C., I.-H. Chung, and J.K. Hollingsworth. *Active Harmony: Towards Automated Performance Tuning*. in *SC'02*. 2002. Baltimore, Maryland.
21. Whaley, R.C. and J.J. Dongarra. *Automatically tuned linear algebra software (ATLAS)*. in *Supercomputing*. 1998. Orlando, FL.
22. Wolf, J. and P.S. Yu, *On Balancing the Load in a Clustered Web Farm*. *ACM Transactions on Internet Technology*, 2001. **1**(2): p. 231-261.