Application of Data Mining to Performance Management of Distributed Enterprise Systems

Sergio Pacheco-Sanchez[1,2],
supervised by {Bryan Scotney[2], Sally McClean[2], and Gerard Parr[2]}

[1] SAP Research CEC Belfast,
Building 1, The Concourse, Queen's Road,
Belfast BT3 9DT, U.K.
sergio.pacheco-sanchez@sap.com

[2] School of Computing and Information Engineering,
University of Ulster, Cromore Road,
Coleraine BT52 1SA, U.K.
{bw.scotney, si.mcclean, gp.parr}@ulster.ac.uk

1 Introduction

The performance evaluation of a computer system involves constructing a suitable model of the system and then using it to predict the system behavior. The model incorporates domain-specific and a priori knowledge about the system inner structure and the workload. The model may be solved either by using analytical techniques or simulation or both. The fundamental parameters needed to specify performance models are arrival times and service demands. The former refers to the frequency of new incoming requests entering the system. The latter indicates the amount of time a request needs to be served, therefore its precise estimation is of vital importance for defining models that are both representative and robust. The number of resource consumption estimation methods that can provide an adequate estimation of the service demand placed on the system [7, 9] is limited compared to the various modeling and evaluation techniques that exist to obtain performance predictions. Essentially, existing techniques for characterization of resource consumption needed for model parameterization are based on direct data measurement and statistical inference. Direct data measurement techniques propose to set up performance probes in the system possibly at different levels, depending on the layers to be analyzed. For example, Magpie [1] automatically extracts system workload characteristics using low-overhead instrumentation. It correlates system generated events to the control flow and to the resource consumption of the requests. Although effective, it is not always possible to adopt instrumentation based approaches of this type in production systems, thus statistical inference is often preferred. Statistical inference encompasses the use of statistics and random sampling to make inferences concerning some unknown aspect of a given population, e.g. response time of service requests. It includes different areas such as predictive inference, estimation theory, among others. It also has different manners of interpreting the concept of probability, namely frequency probability and bayesian probability. For instance, statistical learning techniques are appealing because they may assist in building system models or subsets of them (i.e. workload model) in an automated or semi-automated manner with limited previous knowledge of the system environment. The present work proposes to investigate the way: (i) machine learning can be used in workload characterization tasks; (ii) statistical inference can assist in the correlation of low-level metrics for prediction of system states and in the characterization of resource consumption for system analysis and model parameterization purposes. Initial work done in the latter presented two methods for estimating the service demands of requests based on measurement of their response times [6]. This approach is particularly suitable in scenarios of virtualized systems or for services controlled by third parties. Not surprisingly, the impact of this approach may be of relevance in the near-future, especially when it is forecasted that by 2015 more than 75% of Information Technology infrastructure will be sold as a service by service and infrastructure providers [3]. Preliminary results show the very good accuracy of the proposed approach and a further extension of this investigation is presented in Subsection 3.1. The remainder of the paper is organized as follows. The research challenges of the investigation are presented in the following Subsection. Related work is then introduced in Section 2, followed by a detailed description of the research proposal in Section 3, including (i) and (ii).

1.1 Research Challenges

This subsection offers a short glimpse into the open research fields of computer performance evaluation, namely parameterization of performance models, workload characterization and correlation.

Parameterization of performance models. The challenge of model parameterization is to determine the mean service demand $E[D]$ of the incoming requests at the server such that the response times predicted by the model match accurately those measured in the real system for all possible number of users. The predominant approach to model parameterization makes use of explicit system performance models, such as control-theoretic or queuing-theoretic models. These approaches have achieved major success in many specific performance management applications. However, it is noted that the development of accurate models of complex computing systems is both complicated and highly knowledge-intensive, and moreover, such modeling should become progressively more difficult as systems become increasingly complex and distributed, especially with the advent and mass adoption of Internet component-oriented systems.

Workload characterization. A system workload may be composed of several types of transactions, each one may stress a system in a different manner, therefore possibly consuming system resources differently. Every transaction type is also known as class. The choice of the quantitative parameters (i.e. I/O, CPU time) to be used for determining the number of classes of a workload has to be such that a common description of the composition of each class can be easily obtained. Moreover, the choice of these parameters has to be driven by the intended use of the performance study. Therefore, a key challenge in workload characterization is to determine in an automatic manner the number of workload classes [8]. It is preferable to have a less complex workload and several classes than few classes and a complex workload. A complex workload is composed of classes whose service distribution is difficult to approximate by an exponential. For instance, in presence of high-variability of service demand distributions which are the most challenging to address [7], it seems more practical to decompose the workload in an increased number of classes, such that the service distribution of each class can be approximated by an exponential. This facilitates the modeling of system performance which has become vital for efficiently provisioning services in large distributed systems and for accurately estimating future resource requirements in advance. There is a particular need to investigate inference techniques that assist better in learning mappings between classes and parameters, thus avoiding the use of detailed instrumentation.

Correlation. Another open question in this field is: How do we know what system-level metrics are 'Key Performance Indicators' (KPIs) of a system?. In other words, it is important to understand the relationship between low level metrics and system states (i.e. satisfaction/violation of a quality of service agreement) for achieving system performance management excellence. It is preferred to have small subsets of metrics that capture most of the patterns of performance behavior in a way that is accurate and helps to explain the causes of observed performance effects. Smaller subsets of metrics allow a more efficient representation and evaluation of the model.

2 Related Work

There are two main trends in the literature regarding system modeling for distributed systems. One trend consists of analytical models based on queuing networks [9], control theory [4] and Petri nets [5]. Although mature and mathematically sound, most of the aforementioned models have several limitations: i) the models might be complicated to build, for example, due to absence of knowledge about black-box components; ii) model construction can be prone to human errors; iii) models built are a priori and might not adjust properly to workload or system changes; iv) assumptions that diverge considerably from actual system conditions may have to be made. The second trend is the emerging research rooted in statistical learning techniques using performance data collected through instrumentation. For example, Tree-Augmented Bayesian networks (TANs) are learned from data to correlate low level system measurements, e.g. CPU usage, etc., with system states [2]. Unlike the approaches of the first trend, statistically learned models frequently assume very little or no domain knowledge at all and can be easily updated.

3 Research Proposal

Developing effective models for performance management in distributed computing systems is an important goal of current systems research. Heading in this direction, the present work aims to investigate: (i) mechanisms to improve the process of model parameterization, essential to yield reliable performance predictions; (ii) correlation of low-level metrics for prediction of system states; (iii) automatic or semi-automatic mechanisms for workload characterization. At this stage, preliminary work has been done using a simple *M/M/1* queuing model [6], however, it is needed to further investigate whether other models, either queuing-theoretic or statistical-learning oriented or a combination of both, can capture system performance more accurately. The discovery of hidden associations between low level measurements and the manner by which they impact the final perception of performance may yield in a deeper understanding of the system behavior. Both the mechanisms for model parameterization and workload characterization can benefit from this outcome. For instance, an explicit system performance model can further be refined and its fidelity increased by considering system components not modeled explicitly, but incur in performance degradation. Another example is the investigation of the mapping between a subset of low level measurements that can capture most of the system resource consumption and the various workload classes. This mapping is of high relevance as it can potentially reveal what mix of service demands stress specific system components and ultimately result in a better understanding of the system performance behavior.

3.1 Estimating Service Resource Consumption

In the current work, it is proposed to conduct further model investigation and experimental studies for estimating service resource consumption. This is the extension of previous work [6] that for the first time elaborates the estimation of service demands without knowledge of server utilization measurements. Instead, linear regression (RR) and maximum likelihood (ML) methods are utilized for estimating the request service demands based on the response time measurements only. This means server instrumentation or sampling is not required, therefore the proposed approach significantly simplifies parameterization of performance models when measuring utilization is difficult or unreliable, such as in virtualized systems (because filtering hypervisor system overheads is troublesome and rarely done) or for services controlled by third parties. The ML approach considers in the estimate the entire distribution of the measured response times, thus achieving greater estimation accuracy. The service demands are specified as a function of the number of users N in the model. Furthermore, an intrinsic advantage of the proposed approach is that response times at a server depend on all the latency degradations incurred by requests within the system, therefore they are inherently more comprehensive descriptors of performance as they also account for bandwidth, memory, and I/O contention delays. Very rarely, these components are all directly accounted in models, yet they can be critical performance

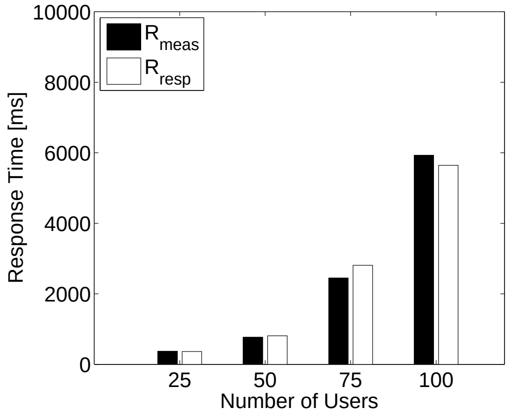

Fig. 1. Comparison of response time predictions.

drivers and are ignored by the utilization approach. During the experimentation phase the RR method service demand estimates (obtained by computing the linear regression of a sequence of response time samples and arrival queue-length averaged on 20 consecutive requests) were used to parameterize a closed *M/M/1* queue model representing an industrial ERP application. Figure 1 shows that the response time predictions obtained through the RR method (R_{resp}) match closely the real response time measurements (R_{meas}) by achieving an accuracy of 90-95%.

The learning from this early work reveals that further sensitivity analyses need to be done by considering various aspects. For example, the impact of non-exponential features in arrivals and service demands for the different request classes and utilization levels is unknown. Regarding the queue-length seen on arrival to the system (simulated with Matlab), it is important to investigate the outcome of using a steady state distribution instead. Another aspect requiring investigation is the utilization of the *M/GI/1/PS* queue. Finally, it is crucial to understand how the lack of information on the residual time of requests impacts on the overall accuracy.

3.2 Statistical Correlation

It is planned to use TAN models to represent the relationship between the low level metrics and the high-level system state as a joint probability distribution. Out of this distribution, a characterization of each metric and its impact on the system state, e.g. violation of a quality of service contract agreed as a maximum two-second response time, can be extracted. The measurements collected from the ERP case study used in subsection 3.1 are vast. It includes several dozen attributes such as sequential read database (DB) rows, number of DB procedure calls, DB update average time/row, time spent in the work process, wait time, CPU time consumed, etc. In fact, there are some metrics with similar names, but different measurement values, therefore it is essential at the same time to investigate the precise differences between them.

3.3 Machine learning

Inspired by earlier work [8], it is planned to investigate what machine learning techniques better assist in characterizing the workload of the ERP application used in the case study of subsection 3.1, possibly using different scenarios to validate the effectiveness of the proposal. Inferring the request classes has direct application to the performance modeling of n-tier on-demand computer systems. These models can benefit from the increased reliability of modeling workloads in separate categories. In fact, this investigation is partially complementary with the planned work described in previous subsection.

4 Acknowledgement

This research has been co-funded by the Northern Ireland government within the InvestNI/SAP program. The author thanks to the MORE project colleagues Giuliano Casale, Stephen Dawson and Stephan Kraft and also to Franklin and Andrea Jones who greatly helped in improving the quality of this paper.

References

1. P. Barham, A. Donnelly, R. Isaacs, and R. Mortier. Using magpie for request extraction and workload modelling. In *OSDI'04*, CA, USA, 2004. USENIX.
2. I. Cohen, M. Goldszmidt, T. Kelly, J. Symons, and J. S. Chase. Correlating instrumentation data to system states: a building block for automated diagnosis and control. In *OSDI'04*, pages 16–16, CA, USA, 2004. USENIX.
3. S. Gaisbauer, J. Kirschnick, N. Edwards, and J. Rolia. Vats: Virtualized-aware automated test service. In *QEST '08*, pages 93–102, 2008.
4. E. Kalyvianaki, T. Charalambous, and S. Hand. Self-adaptive and self-configured cpu resource provisioning for virtualized servers using kalman filters. In *ICAC '09*, pages 117–126. ACM. NY, USA, 2009.
5. S. Kounev and C. Dutz. QPME - A Performance Modeling Tool Based on Queueing Petri Nets. *ACM SIGMETRICS Performance Evaluation Review*, 2009.
6. S. Kraft, S. Pacheco-Sanchez, G. Casale, and S. Dawson. Estimating service resource consumption from response time measurements. In *ValueTools '09*, ICST, Brussels, Belgium, to appear in 2009.
7. J. Rolia and V. Vetland. Correlating resource demand information with arm data for application services. In *WOSP '98*, pages 219–230. ACM. NY, USA, 1998.
8. A. B. Sharma, R. Bhagwan, M. Choudhury, L. Golubchik, R. Govindan, and G. M. Voelker. Automatic request categorization in internet services. In *HotMETRICS'08*. ACM. NY, USA, 2008.
9. Q. Zhang, L. Cherkasova, and E. Smirni. A regression-based analytic model for dynamic resource provisioning of multi-tier applications. In *ICAC '07*, DC, USA, 2007. IEEE.