

PERFORMANCE ENGINEERING OF BUSINESS INFORMATION SYSTEMS

Filling the Gap between High-level Business Services and Low-level Performance Models

Samuel Kounev

Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany

kounev@kit.edu

Keywords: business information systems, performance, scalability, predictive modeling, simulation

Abstract: With the increasing adoption of virtualization and the transition towards Cloud Computing platforms, modern business information systems are becoming increasingly complex and dynamic. This raises the challenge of guaranteeing system performance and scalability while at the same time ensuring efficient resource usage. In this paper, we present a historical perspective on the evolution of model-based performance engineering techniques for business information systems focusing on the major developments over the past several decades that have shaped the field. We survey the state-of-the-art on performance modeling and management approaches discussing the ongoing efforts in the community to increasingly bridge the gap between high-level business services and low level performance models. Finally, we wrap up with an outlook on the emergence of self-aware systems engineering as a new research area at the intersection of several computer science disciplines.

1 INTRODUCTION

Modern business information systems are expected to satisfy increasingly stringent performance and scalability requirements. Most generally, the *performance* of a system refers to the degree to which the system meets its objectives for timeliness and the efficiency with which it achieves this (Smith and Williams, 2002; Kounev, 2008). Timeliness is normally measured in terms of meeting certain response time and/or throughput requirements, *response time* referring to the time required to respond to a user request (e.g., a Web service call or a database transaction), and *throughput* referring to the number of requests or jobs processed per unit of time. Scalability, on the other hand, is understood as the ability of the system to continue to meet its objectives for response time and throughput as the demand for the services it provides increases and resources (typically hardware) are added. Numerous studies, e.g., in the areas of e-business, manufacturing, telecommunications, health care and transportation, have shown that a failure to meet performance requirements can lead to serious financial losses, loss of customers and reputation, and in some cases even to loss of human lives. To avoid

the pitfalls of inadequate Quality-of-Service (QoS), it is important to analyze the expected performance and scalability characteristics of systems during all phases of their life cycle. The methods used to do this are part of the discipline called *Performance Engineering*. Performance Engineering helps to estimate the level of performance a system can achieve and provides recommendations to realize the optimal performance level. The latter is typically done by means of performance models (e.g., analytical queueing models or simulation models) that are used to predict the performance of the system under the expected workload.

However, as systems grow in size and complexity, estimating their performance becomes a more and more challenging task. Modern business information systems based on the Service-Oriented Architecture (SOA) paradigm are often composed of multiple independent *services* each implementing a specific business activity. Services are accessed according to specified workflows representing business processes. Each service is implemented using a set of software components distributed over physical tiers as depicted in Figure 1. Three tiers exist: presentation tier, business logic tier, and data tier. The pre-

sentation tier includes Web servers hosting Web components that implement the presentation logic of the application. The business logic tier normally includes a cluster of application servers hosting business logic components that implement the business logic of the application. Middleware platforms such as Java EE, Microsoft .NET, or Apache Tomcat are often used in this tier to simplify application development by leveraging some common services typically used in enterprise applications. Finally, the data tier includes database servers and legacy systems that provide data management services.

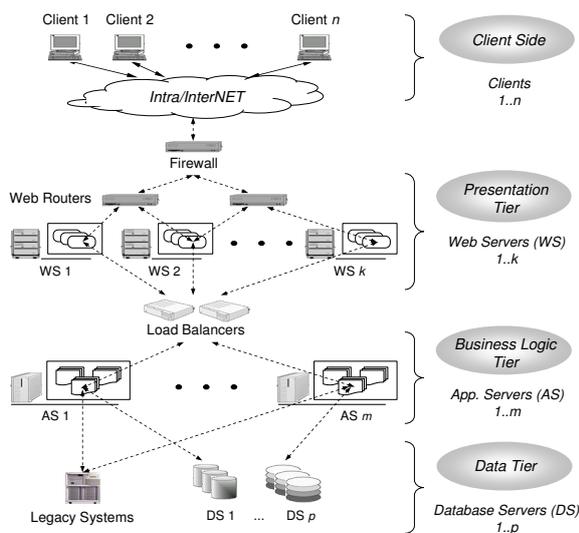


Figure 1: Modern business information system.

The inherent complexity of such architectures makes it difficult to manage their end-to-end performance and scalability. To avoid performance problems, it is essential that systems are subjected to rigorous performance evaluation during the various stages of their lifecycle. At every stage, performance evaluation is conducted with a specific set of goals and constraints. The goals can be classified in the following categories, some of which partially overlap:

Platform selection: Determine which hardware and software platforms would provide the best scalability and cost/performance ratio?

Platform validation: Validate a selected combination of platforms to ensure that taken together they provide adequate performance and scalability.

Evaluation of design alternatives: Evaluate the relative performance, scalability and costs of alternative system designs and architectures.

Performance prediction: Predict the performance of the system for a given workload and configuration scenario.

Performance tuning: Analyze the effect of various deployment settings and tuning parameters on the system performance and find their optimal values.

Performance optimization: Find the components with the largest effect on performance and study the performance gains from optimizing them.

Scalability and bottleneck analysis: Study the performance of the system as the load increases and more hardware is added. Find which system components are most utilized and investigate if they are potential bottlenecks.

Sizing and capacity planning: Determine how much hardware resources are required to guarantee certain performance levels.

Run-time performance and power management: Determine how to vary resource allocations during operation in order to ensure that performance requirements are continuously satisfied while optimizing power consumption in the face of frequent variations in service workloads.

Two broad approaches are used in Performance Engineering for performance evaluation of software systems: performance measurement and performance modeling. In the first approach, load testing tools and benchmarks are used to generate artificial workloads on the system and to measure its performance. In the second approach, performance models are built and then used to analyze the performance and scalability characteristics of the system.

In this paper, we focus on performance modeling since it is normally much cheaper than load testing and has the advantage that it can also be applied in the early stages of system development before the system is available for testing. We present a historical perspective on the evolution of performance modeling techniques for business information systems over the past several decades, focusing on the major developments that have shaped the field, such as the increasing integration of software-related aspects into performance models, the increasing parametrization of models to foster model reuse, the increasing use of automated model-to-model transformations to bridge the gap between models at different levels of abstraction, and finally the increasing use of models at run-time for online performance management.

The paper starts with an overview of classical performance modeling approaches which is followed by an overview of approaches to integrate performance modeling and prediction techniques into the software engineering process. Next, automated model-to-model transformations from architecture-level performance models to classical stochastic performance models are surveyed. Finally, the use of models at

run-time for online performance management is discussed and the paper is wrapped up with some concluding remarks.

2 CLASSICAL PERFORMANCE MODELING

The performance modeling approach to software performance evaluation is based on using mathematical or simulation models to predict the system performance under load. A *performance model* is an abstract representation of the system that relates the workload parameters with the system configuration and captures the main factors that determine the system performance (Menascé et al., 1994). A number of different methods and techniques have been proposed in the literature for modeling software systems and predicting their performance under load. Most of them, however, are based on the same general methodology that proceeds through the steps depicted in Figure 2 (Menascé et al., 2004; Smith and Williams, 2002; Kounev, 2006). First, the goals and objectives of the modeling study are specified. After this, the system is described in detail in terms of its hardware and software architecture. Next, the workload of the system is characterized and a workload model is built. The workload model is used as a basis for building a performance model. Before the model can be used for performance prediction, it has to be validated. This is done by comparing performance metrics predicted by the model with measurements on the real system obtained in a small testing environment. If the predicted values do not match the measured values within an acceptable level of accuracy, then the model must be refined and/or calibrated. Finally, the validated performance model is used to predict the system performance for the deployment configurations and workload scenarios of interest. The model predictions are analyzed and used to address the goals set in the beginning of the modeling study.

Performance models have been employed for performance prediction of software systems since the early seventies. In 1971, Buzen proposed modeling systems using queueing network models and developed solution techniques for several important classes of models. Since then many advances have been made in improving the model expressiveness and developing efficient model analysis techniques as well as accurate approximation techniques. A number of modeling techniques utilizing a range of different performance models have been proposed including standard queueing networks, extended and layered queueing networks, stochastic Petri nets, queueing Petri nets,

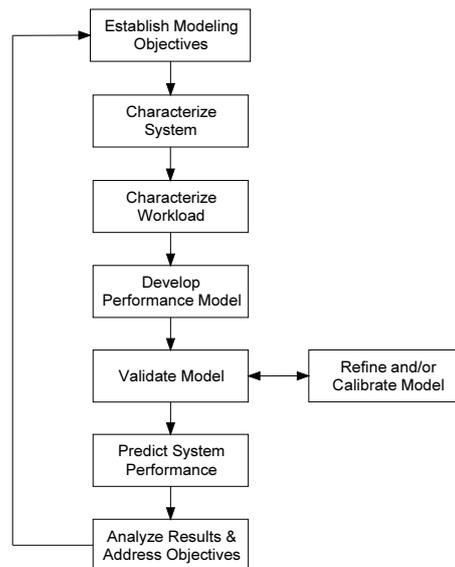


Figure 2: Performance modeling process.

stochastic process algebras, Markov chains, statistical regression models and simulation models. Performance models can be grouped into two main categories: *simulation models* and *analytical models*. One of the greatest challenges in building a good model is to find the right level of abstraction and granularity. A general rule of thumb is: Make the model as simple as possible, but not simpler! Including too much detail might render the model intractable, on the other hand, making it too simple might render it unrepresentative.

2.1 Simulation Models

Simulation models are software programs that mimic the behavior of a system as requests arrive and get processed at the various system resources. Such models are normally stochastic because they have one or more random variables as input (e.g., the times between successive arrivals of requests). The structure of a simulation program is based on the states of the simulated system and events that cause the system state to change. When implemented, simulation programs count events and record the duration of time spent in different states. Based on these data, performance metrics of interest (e.g., the average time a request takes to complete or the average system throughput) can be estimated at the end of the simulation run. Estimates are provided in the form of confidence intervals. A confidence interval is a range with a given probability that the estimated performance metric lies within this range. The main advantage of simulation models is that they are very general and can be made as accurate as desired. However, this accuracy comes at the cost of the time taken to develop

and run the models. Usually, many long runs are required to obtain estimates of needed performance measures with reasonable confidence levels.

Several approaches to developing a simulation model exist. The most time-consuming approach is to use a general purpose programming language such as C++ or Java, possibly augmented by simulation libraries (e.g., CSIM or SimPack, OMNeT++, DESMO-J). Another approach is to use a specialized simulation language such as GPSS/H, Simgen II.5, or MODSIM III. Finally, some simulation packages support graphical languages for defining simulation models (e.g., Arena, Extend, SES/workbench, QPME). A comprehensive treatment of simulation techniques can be found in (Law and Kelton, 2000; Banks et al., 2001).

2.2 Analytical Models

Analytical models are based on mathematical laws and computational algorithms used to derive performance metrics from model parameters. Analytical models are usually less expensive to build and more efficient to analyze compared to simulation models. However, because they are defined at a higher level of abstraction, they are normally less detailed and accurate. Moreover, for models to be mathematically tractable, usually many simplifying assumptions need to be made impairing the model representativeness. Queueing networks and generalized stochastic Petri nets are perhaps the two most popular types of models used in practice.

Queueing networks provide a very powerful mechanism for modeling hardware contention (contention for CPU time, disk access, and other hardware resources). A number of efficient analysis methods have been developed for a class of queueing networks called product-form queueing networks allowing models of realistic size and complexity to be analyzed with a minimum overhead (Bolch et al., 2006). The downside of queueing networks is that they do not provide direct means to model software contention aspects accurately (contention for processes, threads, database connections, and other software resources), as well as blocking, simultaneous resource possession, asynchronous processing, and synchronization aspects. Even though extensions of queueing networks, such as extended queueing networks (MacNair, 1985) and layered queueing networks (also called stochastic rendezvous networks) (Woodside et al., 1995), provide some support for modeling software contention and synchronization aspects, they are often restrictive and inaccurate.

In contrast to queueing networks, generalized

stochastic Petri net models can easily express software contention, simultaneous resource possession, asynchronous processing, and synchronization aspects. Their major disadvantage, however, is that they do not provide any means for direct representation of scheduling strategies. The attempts to eliminate this disadvantage have led to the emergence of queueing Petri nets (Bause, 1993), which combine the modeling power and expressiveness of queueing networks and stochastic Petri nets. Queueing Petri nets enable the integration of hardware and software aspects of system behavior in the same model (Kounev and Buchmann, 2003). A major hurdle to the practical use of queueing Petri nets, however, is that their analysis suffers from the state space explosion problem limiting the size of the models that can be solved. Currently, the only way to circumvent this problem is by using simulation for model analysis (Kounev and Buchmann, 2006).

Details of the various types of analytical models are beyond the scope of this article. The following books can be used as reference for additional information (Bolch et al., 2006; Trivedi, 2002; Bause and Kritzinger, 2002). The Proceedings of the ACM SIGMETRICS Conferences and the Performance Evaluation Journal report recent research results in performance modeling and evaluation. Further relevant information can be found in the Proceedings of the ACM/SPEC International Conference on Performance Engineering (ICPE), the Proceedings of the International Conference on Quantitative Evaluation of Systems (QEST), the Proceedings of the Annual Meeting of the IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), and the Proceedings of the International Conference on Performance Evaluation Methodologies and Tools (VALUETOOLS).

3 SOFTWARE PERFORMANCE ENGINEERING

A major hurdle to the adoption of classical performance modeling approaches in industry is the fact that performance models are expensive to build and require extensive experience and expertise in stochastic modeling which software engineers typically do not possess. To address this issue, over the last fifteen years, a number of approaches have been proposed for integrating performance modeling and prediction techniques into the software engineering process. Efforts were initiated with Smith's seminal work pioneered under the name of *Software Perfor-*

formance Engineering (SPE) (Smith, 1990). Since then a number of languages (i.e., meta-models) for describing performance-relevant aspects of software architectures and execution environments have been developed by the SPE community, the most prominent being the UML SPT profile (UML Profile for Schedulability, Performance and Time) and its successor the UML MARTE profile (UML Profile for Modeling and Analysis of Real-time and Embedded Systems). The latter are extensions of UML (Unified Modeling Language) as the de facto standard modeling language for software architectures. Other proposed architecture-level performance meta-models include SPE-MM (Smith et al., 2005), CSM (Petriu and Woodside, 2007a) and KLAPER (Grassi et al., 2007a). The common goal of these efforts is to enable the automated transformation of architecture-level performance models into analytical or simulation-based performance models that can be solved using classical analysis techniques (see Section 4).

In recent years, with the increasing adoption of Component-Based Software Engineering (CBSE), the SPE community has focused on adapting and extending conventional SPE techniques to support component-based systems. A number of architecture-level performance meta-models for component-based systems have been proposed as surveyed in (Koziolek, 2009). Such meta-models provide means to describe the performance-relevant aspects of software components (e.g., internal control flow and resource demands) while explicitly capturing the influences of their execution context. The idea is that once component models are built they can be reused in multiple applications and execution contexts. The performance of a component-based system can be predicted by means of compositional analysis techniques based on the performance models of its components. Over the last five years, research efforts have been targeted at increasing the level of parametrization of component models to capture additional aspects of their execution context.

An example of a mature modeling language for component-based systems is given by the Palladio Component Model (PCM) (Becker et al., 2009b). In PCM, the component execution context is parameterized to explicitly capture the influence of the component's connections to other components, its allocated hardware and software resources, and its usage profile including service input parameters. Model artifacts are divided among the developer roles involved in the CBSE process, i.e., component developers, system architects, system deployers and domain experts.

4 MODEL-TO-MODEL TRANSFORMATIONS

To bridge the gap between architecture-level performance models and classical stochastic performance models, over the past decade the SPE community has focused on building automated model-to-model transformations which make it possible to exploit existing model solution techniques from the performance evaluation community (Marco and Mirandola, 2006). In the following, we provide an overview of the most common transformations available in the literature.

(Marco and Inverardi, 2004) transform UML models annotated with SPT stereotypes into a multichain queueing network. UML- ψ , the UML Performance Simulator (Marzolla and Balsamo, 2004), transforms a UML instance annotated with the SPT profile to a simulation model. The results from the analysis of the simulation model are reported back to the annotated UML instance (Marco and Mirandola, 2006). Another approach uses the stochastic process algebra PEPA as analysis model (Tribastone and Gilmore, 2008). In this case, only UML activity diagrams are considered, which are annotated with a subset of the MARTE profile. A software tool implementing this method is also available. (Bertolino and Mirandola, 2004) integrate their approach into the Argo-UML modeling tool, using the RT-UML performance annotation profile. An execution graph and a queueing network serve as the target analysis formalisms.

Other approaches use UML, but do not use standardized performance profile annotations: (Gu and Petriu, 2002) use XSLT, the eXtensible Stylesheet Language Transformations, to execute a graph pattern based transformation from a UML instance to LQNs. Instead of annotating the UML model, it has to be modeled in a way so that the transformation can identify the correct patterns in the model. (Bernardi et al., 2002) consider only UML statecharts and sequence diagrams. A transformation written in Java turns the model into GSPN sub-models that are then combined into a final GSPN. (Gomaa and Menascé, 2001) use UML with custom XML performance annotation. The performance model is not described in detail, but appears to be based on queueing networks. (Wu and Woodside, 2004) use UML component models together with a custom XML component performance specification language. LQN solvers are used for the analysis.

Further approaches exist that are not based on UML: (Bondarev et al., 2004; Bondarev et al., 2005) build on the ROBOCOP component model and use proprietary simulation framework for model analysis.

(Eskenazi et al., 2004) propose a custom control flow graph model notation and custom simulation framework. (Hissam et al., 2002) employ the COMTEK component technology, coupled with a proprietary analysis framework. (Sitaraman et al., 2001) specify component composition and performance characteristics using a variant of the big-O notation. The runtime analysis is not discussed in detail.

Several model-to-model transformations have been developed for the Palladio Component Model (PCM). Two solvers are based on a transformation to Layered Queueing Networks (LQNs) (Koziolek and Reussner, 2008) and a transformation to Stochastic Regular Expressions (Koziolek, 2008), respectively. Stochastic Regular Expressions can be solved analytically with very low overhead, however, they only support single user scenarios. (Henss, 2010) proposes a PCM transformation to OMNeT++, focusing on a realistic network infrastructure closer to the OSI reference network model. The PCM-Bench tool comes with the SimuCom simulator (Becker et al., 2009a) which is based on a model-to-text transformation used to generate Java code that builds on DESMO-J, a general-purpose simulation framework. The code is then compiled on-the-fly and executed. SimuCom is tailored to support all of the PCM features directly and covers the whole PCM meta-model.

Finally, a number of intermediate languages (or kernel languages) for specifying software performance information have been proposed in the literature. The aim of such efforts is to reduce the overhead for building transformations, i.e., only $M + N$ instead of $M \cdot N$ transformations have to be developed for M source and N target meta-models (Marco and Mirandola, 2006). Some examples of intermediate languages include SPE-MM (Smith et al., 2005), KLAPER (Kernel LAnguage for PErformance and Reliability analysis) (Grassi et al., 2007b) and CSM (Core Scenario Model) (Petriu and Woodside, 2007b).

5 RUN-TIME PERFORMANCE MANAGEMENT

With the increasing adoption of virtualization and the transition towards Cloud Computing platforms, modern business information systems are becoming increasingly complex and dynamic. The increased complexity is caused by the introduction of virtual resources and the resulting gap between logical and physical resource allocations. The increased dynamism is caused by the complex interactions between the applications and services sharing the physical in-

frastructure. In a virtualized service-oriented environment changes are common during operation, e.g., new services and applications can be deployed on-the-fly, service workflows and business processes can be modified dynamically, hardware resources can be added and removed from the system, virtual machines (VMs) can be migrated between servers, resources allocated to VMs can be modified to reflect changes in service workloads and usage profiles. To ensure adequate performance and efficient resource utilization in such environments, capacity planning needs to be done on a regular basis during operation. This calls for *online* performance prediction mechanisms.

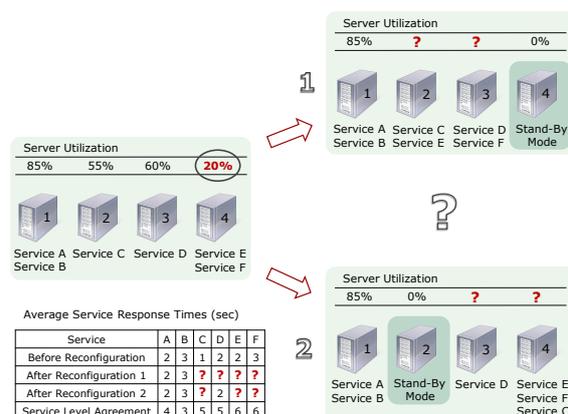


Figure 3: Online performance prediction scenario.

An example of a scenario where online performance prediction is needed is depicted in Figure 3. A service-oriented system made of four servers hosting six different services is shown including information on the average service response times, the response time service level agreements (SLAs) and the server utilization. Now assume that due to a change in the demand for services E and F, the average utilization of the fourth server has dropped down to 20% over an extended period of time. To improve the system's efficiency, it is considered to switch one of the servers to stand-by mode after migrating its services to other servers. Two possible ways to reconfigure the system are shown. To ensure that reconfiguring the system would not break the SLAs, the system needs a mechanism to predict the effect of the reconfiguration on the service response times.

Given the variety of changes that occur in modern service-oriented environments, online performance prediction techniques must support variations at all levels of the system including variations in service workloads and usage profiles, variations in the system architecture, as well as variations in the deployment and execution environment (virtualization, middleware, etc). To predict the impact of such variations,

architecture-level performance models are needed at run-time that explicitly capture the influences of the system architecture, its configuration, and its workload and usage profiles.

While, as discussed in the previous two sections, many architecture-level performance prediction techniques exist in the literature, most of them suffer from two significant drawbacks which render them impractical for use at run-time: i) performance models provide limited support for reusability and customization, ii) performance models are static and maintaining them manually during operation is prohibitively expensive (Woodside et al., 2007).

While techniques for component-based performance engineering have contributed a lot to facilitate model reusability, there is still much work to be done on further parameterizing component models before they can be used for online performance prediction. In particular, current techniques do not provide means to model the layers of the component execution environment explicitly. The performance influences of the individual layers, the dependencies among them, as well as the resource allocations at each layer should be captured as part of the models. This is necessary in order to be able to predict at run-time how a change in the execution environment (e.g., modifying resource allocations at the virtualization layer) would affect the overall system performance.

As to the second issue indicated above, the heart of the problem is in the fact that architecture-level performance models are normally designed for offline use and as such they are decoupled from the system components they represent. Thus, models do not capture dynamic aspects of the environment and therefore they need to be updated manually after every change in the system configuration or workload. Given the frequency of such changes, the amount of effort involved in maintaining performance models is prohibitive and therefore in practice such models are rarely used after deployment. Even though some techniques have been proposed to automatically construct stochastic performance models at run-time, e.g., (Menascé et al., 2007; Mos, 2004), such techniques abstract the system at a very high level without taking into account its architecture and configuration.

To address the challenges described above, current research efforts are focusing on developing *online* architecture-level performance models designed specifically for use at run-time. Such models aim at capturing all information, both static and dynamic, relevant to predicting the system's performance on-the-fly. They are intended to be integrated into the system components and to be maintained and updated automatically by the underlying execution platform

(virtualization and middleware) reflecting the evolving system environment.

Online performance models will make it possible to answer performance-related queries that arise during operation such as: What would be the effect on the performance of running applications if a new application is deployed in the virtualized infrastructure or an existing application is migrated from one physical server to another? How much resources need to be allocated to a newly deployed application to ensure that SLAs are satisfied? How should the system configuration be adapted to avoid performance issues or inefficient resource usage arising from changing customer workloads?

The ability to answer queries such as the above provides the basis for implementing techniques for *self-aware* performance and resource management (Kounev et al., 2010). Such techniques will be triggered automatically during operation in response to observed or forecast changes in application workloads. The goal will be to *proactively* adapt the system to such changes in order to avoid anticipated QoS problems or inefficient resource usage. The adaptation will be performed in an autonomic fashion by considering a set of possible system reconfiguration scenarios (e.g. changing VM placement and/or resource allocations) and exploiting the online performance models to predict the effect of such reconfigurations before making a decision.

Self-aware systems engineering is currently emerging as a new research area at the intersection of several computer science disciplines including software architecture, computer systems modeling, autonomic computing, distributed systems, and more recently, Cloud Computing and Green IT (Descartes Research Group, 2011; Kounev, 2011). It raises a number of big challenges that represent emerging hot topics in the systems engineering community and will be subject of long-term fundamental research in the years to come. The resolution of these challenges promises to revolutionize the field of systems engineering by enabling guaranteed QoS, lower operating costs and improved energy efficiency.

6 CONCLUDING REMARKS

We presented a historical perspective on the evolution of model-based performance engineering techniques for business information systems, focusing on the major developments over the past four decades that have shaped the field, such as the increasing integration of software-related aspects into performance models, the increasing parametrization of models to

foster model reuse, the increasing use of automated model-to-model transformations to bridge the gap between models at different levels of abstraction, and finally the increasing use of models at run-time for online performance management. We surveyed the state-of-the-art on performance modeling and management approaches discussing the ongoing efforts in the community to increasingly bridge the gap between high-level business services and low level performance models. Finally, we concluded with an outlook on the emergence of self-aware systems engineering as a new research area at the intersection of several computer science disciplines.

REFERENCES

- Banks, J., Carson, J. S., Nelson, B. L., and Nicol, D. M. (2001). *Discrete-Event System Simulation*. Prentice Hall, Upper Saddle River, NJ 07458, third edition.
- Bause, F. (1993). Queuing Petri Nets - A formalism for the combined qualitative and quantitative analysis of systems. In *Proceedings of the 5th International Workshop on Petri Nets and Performance Models, Toulouse, France, October 19-22*.
- Bause, F. and Kritzinger, F. (2002). *Stochastic Petri Nets - An Introduction to the Theory*. Vieweg Verlag, second edition.
- Becker, S., Koziolok, H., and Reussner, R. (2009a). The palladio component model for model-driven performance prediction. *J. Syst. Softw.*, 82(1):3–22.
- Becker, S., Koziolok, H., and Reussner, R. (2009b). The Palladio component model for model-driven performance prediction. *Journal of Syst. and Softw.*, 82:3–22.
- Bernardi, S., Donatelli, S., and Merseguer, J. (2002). From UML sequence diagrams and statecharts to analysable petri net models. In *Proc. on WOSP '02*, pages 35–45.
- Bertolino, A. and Mirandola, R. (2004). CB-SPE tool: Putting component-based performance engineering into practice. In *Component-Based Software Engineering*, volume 3054 of *LNCS*, pages 233–248.
- Bolch, G., Greiner, S., Meer, H. D., and Trivedi, K. S. (2006). *Queueing Networks and Markov Chains: Modeling and Performance Evaluation with Computer Science Applications*. John Wiley & Sons, Inc., 2nd edition.
- Bondarev, E., de With, P., Chaudron, M., and Muskens, J. (2005). Modelling of input-parameter dependency for performance predictions of component-based embedded systems. In *Proc. on EUROMICRO '05*, pages 36–43.
- Bondarev, E., Muskens, J., de With, P., Chaudron, M., and Lukkien, J. (2004). Predicting real-time properties of component assemblies: a scenario-simulation approach. pages 40–47.
- Descartes Research Group (2011). <http://www.descartes-research.net>.
- Eskenazi, E., Fioukov, A., and Hammer, D. (2004). Performance prediction for component compositions. In *Component-Based Software Engineering*, volume 3054 of *Lecture Notes in Computer Science*, pages 280–293.
- Gomaa, H. and Menascé, D. (2001). Performance engineering of component-based distributed software systems. In *Performance Engineering*, volume 2047 of *LNCS*, pages 40–55.
- Grassi, V., Mirandola, R., and Sabetta, A. (2007a). Filling the gap between design and performance/reliability models of component-based systems: A model-driven approach. *Journal of Systems and Software*, 80(4):528–558.
- Grassi, V., Mirandola, R., and Sabetta, A. (2007b). Filling the gap between design and performance/reliability models of component-based systems: A model-driven approach. *J. Syst. Softw.*, 80(4):528–558.
- Gu, G. P. and Petriu, D. C. (2002). XSLT transformation from UML models to LQN performance models. In *Proc. on WOSP '02*, pages 227–234.
- Henss, J. (2010). Performance prediction for highly distributed systems. In *Proc. on WCOP '10*, volume 2010-14, pages 39–46. Karlsruhe Institute of Technology.
- Hissam, S., Moreno, G., Stafford, J., and Wallnau, K. (2002). Packaging predictable assembly. In *Component Deployment*, volume 2370 of *Lecture Notes in Computer Science*, pages 108–124.
- Kounev, S. (2006). Performance Modeling and Evaluation of Distributed Component-Based Systems using Queuing Petri Nets. *IEEE Transactions on Software Engineering*, 32(7):486–502.
- Kounev, S. (2008). *Wiley Encyclopedia of Computer Science and Engineering*, chapter Software Performance Evaluation. Wiley-Interscience, John Wiley & Sons Inc, ISBN-10: 0471383937, ISBN-13: 978-0471383932.
- Kounev, S. (2011). Self-Aware Software and Systems Engineering: A Vision and Research Roadmap. In *Proceedings of Software Engineering 2011 (SE2011), Nachwuchswissenschaftler-Symposium*.
- Kounev, S., Brosig, F., Huber, N., and Reussner, R. (2010). Towards self-aware performance and resource management in modern service-oriented systems. In *Proc. of the 7th IEEE Intl. Conf. on Services Computing (SCC 2010)*. IEEE Computer Society.
- Kounev, S. and Buchmann, A. (2003). Performance Modelling of Distributed E-Business Applications using Queuing Petri Nets. In *Proceedings of the 2003 IEEE International Symposium on Performance Analysis of Systems and Software - ISPASS2003, Austin, Texas, USA, March 20-22*.
- Kounev, S. and Buchmann, A. (2006). SimQPN - a tool and methodology for analyzing queuing Petri net models by means of simulation. *Performance Evaluation*, 63(4-5):364–394.

- Koziolok, H. (2008). *Parameter dependencies for reusable performance specifications of software components*. PhD thesis, University of Karlsruhe (TH).
- Koziolok, H. (2009). Performance evaluation of component-based software systems: A survey. *Performance Evaluation*.
- Koziolok, H. and Reussner, R. (2008). A model transformation from the palladio component model to layered queueing networks. In *Proc. on SIPEW '08*, pages 58–78.
- Law, A. and Kelton, D. W. (2000). *Simulation Modeling and Analysis*. Mc Graw Hill Companies, Inc., third edition.
- MacNair, E. A. (1985). An introduction to the Research Queueing Package. In *WSC '85: Proceedings of the 17th conference on Winter simulation*, pages 257–262, New York, NY, USA. ACM Press.
- Marco, A. D. and Inverardi, P. (2004). Compositional generation of software architecture performance QN models. *Software Architecture, Working IEEE/IFIP Conf. on*, 0:37.
- Marco, A. D. and Mirandola, R. (2006). Model transformation in software performance engineering. In *QoSA*.
- Marzolla, M. and Balsamo, S. (2004). UML-PSI: The UML performance simulator. *Quantitative Eval. of Syst.*, 0:340–341.
- Menascé, D., Ruan, H., and Goma, H. (2007). QoS management in service-oriented architectures. *Performance Evaluation*, 64(7-8):646–663.
- Menascé, D. A., Almeida, V., and Dowdy, L. W. (1994). *Capacity Planning and Performance Modeling - From Mainframes to Client-Server Systems*. Prentice Hall, Englewood Cliffs, NG.
- Menascé, D. A., Almeida, V., and Dowdy, L. W. (2004). *Performance by Design*. Prentice Hall.
- Mos, A. (2004). *A Framework for Adaptive Monitoring and Performance Management of Component-Based Enterprise Applications*. PhD thesis, Dublin City University, Ireland.
- Petriu, D. and Woodside, M. (2007a). An intermediate metamodel with scenarios and resources for generating performance models from UML designs. *Software and Systems Modeling (SoSyM)*, 6(2):163–184.
- Petriu, D. and Woodside, M. (2007b). An intermediate metamodel with scenarios and resources for generating performance models from UML designs. *Softw. and Syst. Modeling*, 6(2):163–184.
- Sitaraman, M., Kulczycki, G., Krone, J., Ogden, W. F., and Reddy, A. L. N. (2001). Performance specification of software components. *SIGSOFT Softw. Eng. Notes*, 26(3):3–10.
- Smith, C. U. (1990). *Performance Engineering of Software Systems*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.
- Smith, C. U., Llad, C. M., Cortellessa, V., Di Marco, A., and Williams, L. G. (2005). From UML models to software performance results: an SPE process based on XML interchange formats. In *WOSP '05: Proceedings of the 5th International Workshop on Software and Performance*, pages 87–98, New York, NY, USA. ACM Press.
- Smith, C. U. and Williams, L. G. (2002). *Performance Solutions - A Practical Guide to Creating Responsive, Scalable Software*. Addison-Wesley.
- Tribastone, M. and Gilmore, S. (2008). Automatic extraction of PEPA performance models from UML activity diagrams annotated with the MARTE profile. In *Proc. on WOSP '08*.
- Trivedi, K. S. (2002). *Probability and Statistics with Reliability, Queueing and Computer Science Applications*. John Wiley & Sons, Inc., second edition.
- Woodside, M., Franks, G., and Petriu, D. (2007). The future of software performance engineering. In *Future of Software Engineering (FOSE'07)*, pages 171–187, Los Alamitos, CA, USA. IEEE Computer Society.
- Woodside, M., Neilson, J., Petriu, D., and Majumdar, S. (1995). The Stochastic Rendezvous Network Model for Performance of Synchronous Client-Server-Like Distributed Software. *IEEE Transactions on Computers*, 44(1):20–34.
- Wu, X. and Woodside, M. (2004). Performance modeling from software components. *SIGSOFT Softw. Eng. Notes*, 29(1):290–301.