Reactive vs. Proactive Detection of Quality of Service Problems

Lars Grunske[a], Ayman Amin[b]

[a]Institute of Software Technology, University of Stuttgart, Germany
[b]Faculty of ICT, Swinburne University of Technology, Australia

Abstract: This paper summarizes our earlier contributions on reactive and proactive detection of quality of service problems. The first contribution is applying statistical control charts to reactively detect QoS violations. The second contribution is applying time series modeling to proactively detect potential QoS violations.

1 Introduction

Software systems may suffer at runtime from changes in their operational environment or/and requirements specification, so they need to be adapted to satisfy the changed environment or/and specifications [CGK+11]. The research community has developed a number of approaches to building adaptive systems that respond to these changes such as Rainbow [GCH+04]. Currently, several approaches have been proposed to monitor QoS attributes at runtime with the goal of reactively detecting QoS violations (e.g. [MP11]). We will present our reactive and proactive techniques in the following.

2 Reactive detection of QoS violations based on control charts

The reactive approaches detect QoS violations by observing the running system and determining QoS values [MP11] and checking if they exceed a predefined threshold. The main limitation of these approaches is that they do not have statistical confidence in detecting QoS violations. To address this limitation, we propose a statistical approach [ACG11, ACG12b] based on control charts for the runtime detection of QoS violations. This approach consists of four phases: (1) Estimating the running software system capability (current normal behavior) in terms of descriptive statistics, i.e. mean, std, and confidence interval of the QoS attributes; (2) Building a control chart (esp. CUSUM) using given QoS requirements; (3) After each new QoS observation, updating the chart statistic and checking for statistically significant violations; (4) In case of detecting violations, providing warning signals.

3 Proactive Detection of QoS violations based on time series modeling

Predicting future values of Quality of Service (QoS) attributes can assist in the control of software intensive systems by preventing QoS violations before they happen. Currently, many approaches prefer ARIMA models for this task, and assume the QoS attributes' behavior can be linearly modeled. However, our analysis of real QoS datasets shows that they are characterized by a highly dynamic and mostly nonlinear behavior with volatility clustering (time-varying variation) to the extent that existing ARIMA models cannot guarantee accurate QoS forecasting, which can introduce crucial problems such as proactively triggering unrequired adaptations and thus leading to follow-up failures and increased costs. To address this limitation, we propose two automated forecasting approaches based on time series modeling. The first forecasting approach [AGC12] addresses the nonlinearity characteristic of QoS values. This forecasting approach integrates linear and nonlinear time series models [BJ76] and automatically, without human intervention, selects and constructs the best suitable forecasting model to fit the QoS attributes' dynamic behavior and provide accurate forecasting for QoS measures and violations. The second forecasting approach [ACG12a] addresses the QoS volatility by exploiting the ability of generalized autoregressive conditional heteroscedastic (GARCH) models to model the high volatility [Eng82]. This approach basically integrates ARIMA and GARCH models to capture the QoS volatility and provide accurate forecasting for QoS measures and violations Using real-world QoS datasets of Web services we evaluate the accuracy and performance aspects of the proposed forecasting approaches.

References

[ACG11] Ayman Amin, Alan Colman, and Lars Grunske. Using Automated Control Charts for the Runtime Evaluation of QoS Attributes. In *Proc. of the 13ht IEEE Int. High Assurance Systems Engineering Symposium*, pages 299–306. IEEE Computer Society, 2011.

[ACG12a] Ayman Amin, Alan Colman, and Lars Grunske. An Approach to Forecasting QoS Attributes of Web Services Based on ARIMA and GARCH Models. In *Proc. of the 19th Int. Conf. on Web Services*, pages 74–81. IEEE, 2012.

[ACG12b] Ayman Amin, Alan Colman, and Lars Grunske. Statistical Detection of QoS Violations Based on CUSUM Control Charts. In *Proc. of the 3rd ACM/SPEC Int. Conf. on Performance Engineering*, pages 97–108. ACM, 2012.

[AGC12] Ayman Amin, Lars Grunske, and Alan Colman. An automated approach to forecasting QoS attributes based on linear and non-linear time series modeling. In *Proc. of the 27th IEEE/ACM Int. Conf. on Automated Software Engineering*, pages 130–139. IEEE, 2012.

[BJ76] George E. P. Box and Gwilym M. Jenkins. *Time Series Analysis: Forecasting and Control*. HoldenDay, San Francisco, 1976.

[CGK+11] Radu Calinescu, Lars Grunske, Marta Z. Kwiatkowska, Raffaela Mirandola, and Giordano Tamburrelli. Dynamic QoS Management and Optimization in Service-Based Systems. *IEEE Trans. Software Eng.*, 37(3):387–409, 2011.

[Eng82] R.F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, pages 987–1007, 1982.

[GCH+04] David Garlan, Shang-Wen Cheng, An-Cheng Huang, Bradley Schmerl, and Peter Steenkiste. Rainbow: Architecture-based self-adaptation with reusable infrastructure. *Computer*, 37(10):46–54, 2004.

[MP11] Raffaela Mirandola and Pasqualina Potena. A QoS-based framework for the adaptation of service-based systems. *Scalable Computing: Practice and Experience*, 12(1), 2011.