

## Chapter 9

# Data Mining System Execution Traces to Validate Distributed System Quality-of-Service Properties

**James H. Hill**

*Indiana University-Purdue University Indianapolis, USA*

### ABSTRACT

*System Execution Modeling (SEM) tools enable distributed system testers to validate Quality-of-Service (QoS) properties, such as end-to-end response time, throughput, and scalability, during early phases of the software lifecycle. Analytical capabilities of QoS properties, however, are traditionally bounded by a SEM tool's capabilities. This chapter discusses how to mine system execution traces, which are a collection of log messages describing events and states of a distributed system throughout its execution lifetime, generated by distributed systems so that the validation of QoS properties is not dependent on a SEM tool's capabilities. The author uses a real-life case study to illustrate how data mining system execution traces can assist in discovering potential performance bottlenecks using system execution traces.*

### INTRODUCTION

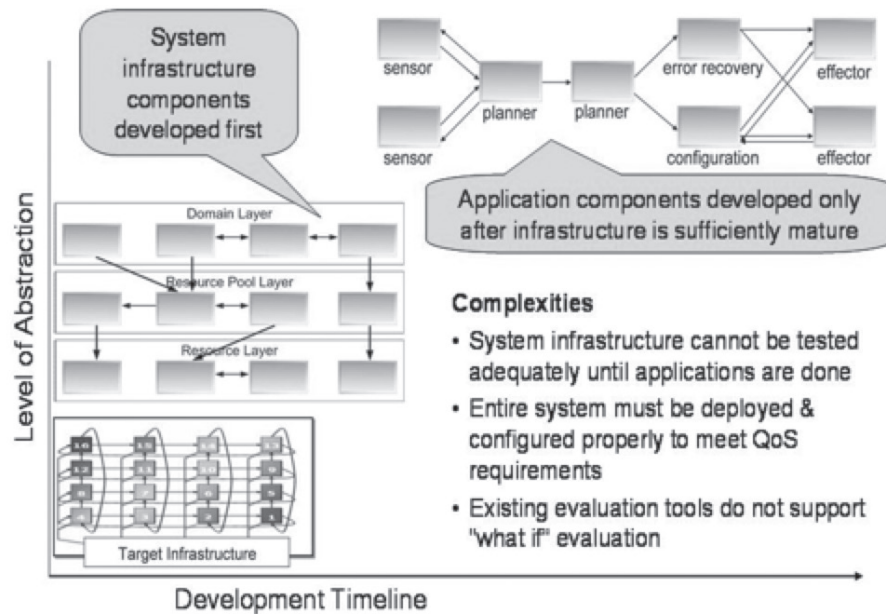
Enterprise distributed systems, such as mission avionic systems, traffic management systems, and shipboard computing environments, are transitioning to next-generation middleware, such as service-oriented middleware (Pezzini & Natis, 2007) and component-based software engineering (Heineman & Council, 2001). Although next-generation middleware is improving enterprise distributed system functional properties (i.e.,

its operational scenarios), Quality-of-Service (QoS) properties (e.g., end-to-end response time, throughput, and scalability) are not validated until late in the software lifecycle, i.e., during system integration time. This is due in part to the *serialized-phasing development* problem (Rittel & Webber, 1973).

As illustrated in Figure 1, in serialized-phasing development, the infrastructure and application-level system entities, such as components that encapsulate common services, are developed during different phases of the software lifecycle. Software design decisions that affect QoS properties,

DOI: 10.4018/978-1-60960-067-9.ch009

Figure 1. Overview of serialized-phased development in distributed systems



however, are typically not discovered until final stages of development, e.g., at system integration time, which is too late in the software lifecycle to resolve performance bottlenecks in an efficient and cost effective manner (Mann, 1996; Snow & Keil, 2001; Woodside, Franks, & Petriu, 2007).

System Execution Modeling (SEM) tools (Smith & Williams, 2001), which are a form of model-driven engineering (Schmidt, 2006), assist distributed system developers in overcoming the serialized-phasing development problem shown in Figure 1. SEM tools use domain-specific modeling languages (Ledeczi, Maroti, Karsai, & Nordstrom, 1999) to capture both platform-independent attributes (such as structural and behavioral concerns of the system) and platform-specific attributes (such as the target architecture of the system) as high-level models. Model interpreters then transform constructed models into source code for the target architecture. This enables distributed system testers to validate QoS properties continuously throughout the software lifecycle while the “real” system is still under development. Likewise, as development of the real

system is complete, distributed system testers can incrementally replace *faux* portions of the system with its real counterpart to produce more realistic QoS validation results.

Although SEM tools enable distributed system developers and testers to validate distributed system QoS properties during early phases of the software lifecycle, QoS validation capabilities are typically bounded to a SEM tool’s analytical capabilities. In order to validate QoS properties unknown to a SEM tool, distributed system testers have the following options:

- **Use handcrafted solutions:** This option typically occurs outside of the SEM. Moreover, this option is traditionally not applicable across different application domains because it is an *ad hoc* solution, e.g., handcrafting a solution to validate event response-time based on priority in a proprietary system;
- **Leverage model transformations to convert models to a different SEM tool:** This option implies the source and target

22 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:  
[www.igi-global.com/chapter/data-mining-system-execution-traces/46896](http://www.igi-global.com/chapter/data-mining-system-execution-traces/46896)

## Related Content

---

### Kernel Width Selection for SVM Classification: A Meta-Learning Approach

Ali Smith and Kate A. Smith (2008). *Mathematical Methods for Knowledge Discovery and Data Mining* (pp. 101-115).

[www.irma-international.org/chapter/kernel-width-selection-svm-classification/26135](http://www.irma-international.org/chapter/kernel-width-selection-svm-classification/26135)

### Technology Transfer Projects at the University-Industry Interface: A Case Study Analysis From the UK

Martin George Wynn (2020). *Current Issues and Trends in Knowledge Management, Discovery, and Transfer* (pp. 272-304).

[www.irma-international.org/chapter/technology-transfer-projects-at-the-university-industry-interface/244888](http://www.irma-international.org/chapter/technology-transfer-projects-at-the-university-industry-interface/244888)

### Application of Data Mining Techniques in Weather Forecasting

ThippaReddy Gadekallu, Bushra Kidwai, Saksham Sharma, Rishabh Pareek and Sudheer Karnam (2019). *Sentiment Analysis and Knowledge Discovery in Contemporary Business* (pp. 162-174).

[www.irma-international.org/chapter/application-of-data-mining-techniques-in-weather-forecasting/210969](http://www.irma-international.org/chapter/application-of-data-mining-techniques-in-weather-forecasting/210969)

### Applications of Domain-Specific Predictive Analytics Applied to Big Data

Ravi Kumar Poluru, S. Bharath Bhushan, Basha Syed Muzamil, Praveen Kumar Rayani and Praveen Kumar Reddy (2019). *Sentiment Analysis and Knowledge Discovery in Contemporary Business* (pp. 289-306).

[www.irma-international.org/chapter/applications-of-domain-specific-predictive-analytics-applied-to-big-data/210976](http://www.irma-international.org/chapter/applications-of-domain-specific-predictive-analytics-applied-to-big-data/210976)

### Mobile Marketing: The Challenges of the New Direct Marketing Channel and the Need for Automatic Targeting and Optimization Tools

Giovanni Giuffrida, Diego Reforgiato, Catarina Sismeiro and Giuseppe Tribulato (2011). *Knowledge Discovery Practices and Emerging Applications of Data Mining: Trends and New Domains* (pp. 151-173).

[www.irma-international.org/chapter/mobile-marketing-challenges-new-direct/46895](http://www.irma-international.org/chapter/mobile-marketing-challenges-new-direct/46895)