A simulation process for asynchronous event processing systems: Evaluating performance and availability in transaction models

Anakreon Mentis, Panagiotis Katsaros *, Lefteris Angelis

Department of Informatics, Aristotle University of Thessaloniki, 54124 Thessaloniki, Greece

ABSTRACT

Simulation is essential for understanding the performance and availability behavior of complex systems, but there are significant difficulties when trying to simulate systems with multiple components, which interact with asynchronous communication. A systematic process is needed, in order to cope with the complexity of asynchronous event processing and the failure semantics of the interacting components. We address this problem by introducing an approach that combines *formal techniques* for faithful representation of the complex system effects and a *statistical analysis* for simultaneously studying multiple simulation outcomes, in order to interpret them. Our process has been successfully applied to a synthetic workload for distributed transaction processing. We outline the steps followed towards generating a credible simulation model and subsequently we report and interpret the results of the applied statistical analysis. This serves as a proof of concept that the proposed simulation process can be also effective in other asynchronous system contexts, like for example distributed group communication systems, file systems and so on.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Systems with asynchronous event processing involve components that communicate by exchanging messages, which are not delivered in the order that are sent or are even lost. Examples include distributed systems for transaction processing, file systems, data services, group communication services, etc. There are two sources of complexity in the specification, design and implementation of these systems: (i) components can be in one of a multitude of states, when they receive a message and (ii) systems involve some failure semantics. Simulation-based analysis for performance and availability must rely on *verifiably correct models*, as well as on *simulation output analyses* that highlight the main factors that determine whether the anticipated goals can be met.

This work introduces *a coherent simulation process that simplifies the development and analysis of complex models for asynchronous event processing systems*. The process steps include automata-based model specification, model checking, simulation code generation, model validation and simulation output statistical analysis. The whole process has been successfully applied for evaluating the effects of transaction models and their associated parameters on system performance and availability.

Analysis of transaction processing models is supported by the ACID Model Checker [14] and the ACID Sim Tools [16], developed by the authors. Our tool support provides a library of methods for simulating basic operations, such as storing a transaction entry in a log file or attempting a lock on an object. A state machine specification of the system model is checked for the expected correctness properties. The verified model is then used to generate simulation source code that

---

* Corresponding author. Tel.: +30 2310 998532; fax: +30 2310 998419.
  E-mail addresses: anakreon@csd.auth.gr (A. Mentis), katsaros@csd.auth.gr (P. Katsaros), lef@csd.auth.gr (L. Angelis).

invokes the library methods and manages the interactions between model entities. The derived formally-validated simulator is subsequently used in simulations of key execution scenarios (e.g. local to highly distributed transactions, I/O bound or CPU bound resource contention) to find out if the model achieves the expected performance and availability goals. An appropriate statistical analysis explores the feasibility of the performance and availability goals.

The first process steps have been also employed in [18], a related work for a different statistical analysis that specifically focuses on the evaluation of complex metric interactions. In the present article we introduce a complete simulation process that aims to determine whether the anticipated performance and availability goals can be met, rather than a statistical analysis to address existing metric correlations.

Section 2 introduces the steps of the proposed simulation process. Section 3 presents how the process is applied to transaction processing models, as a proof of concept for its effectiveness on complex system models. A synthetic transaction workload scenario is provided to demonstrate the application of the process steps. Section 4 focuses on the statistical analysis of the obtained simulation results towards evaluating the feasibility of expected performance and availability goals. Results for the transaction processing problem of Section 3 are provided. Related work is presented in Section 5 and the paper concludes with remarks on the effectiveness of our proposal and future research directions.

2. Model development and simulation process

Fig. 1 presents the simulation process for an asynchronous system model to evaluate the performance and the availability of a distributed system. The steps of the process are:

1. *Model definition*, by means of state machines specifying the model entities and their interactions. As an example, the classical 2PC transaction model [3] can be defined in terms of states and transitions for the transaction coordinator and worker entities.
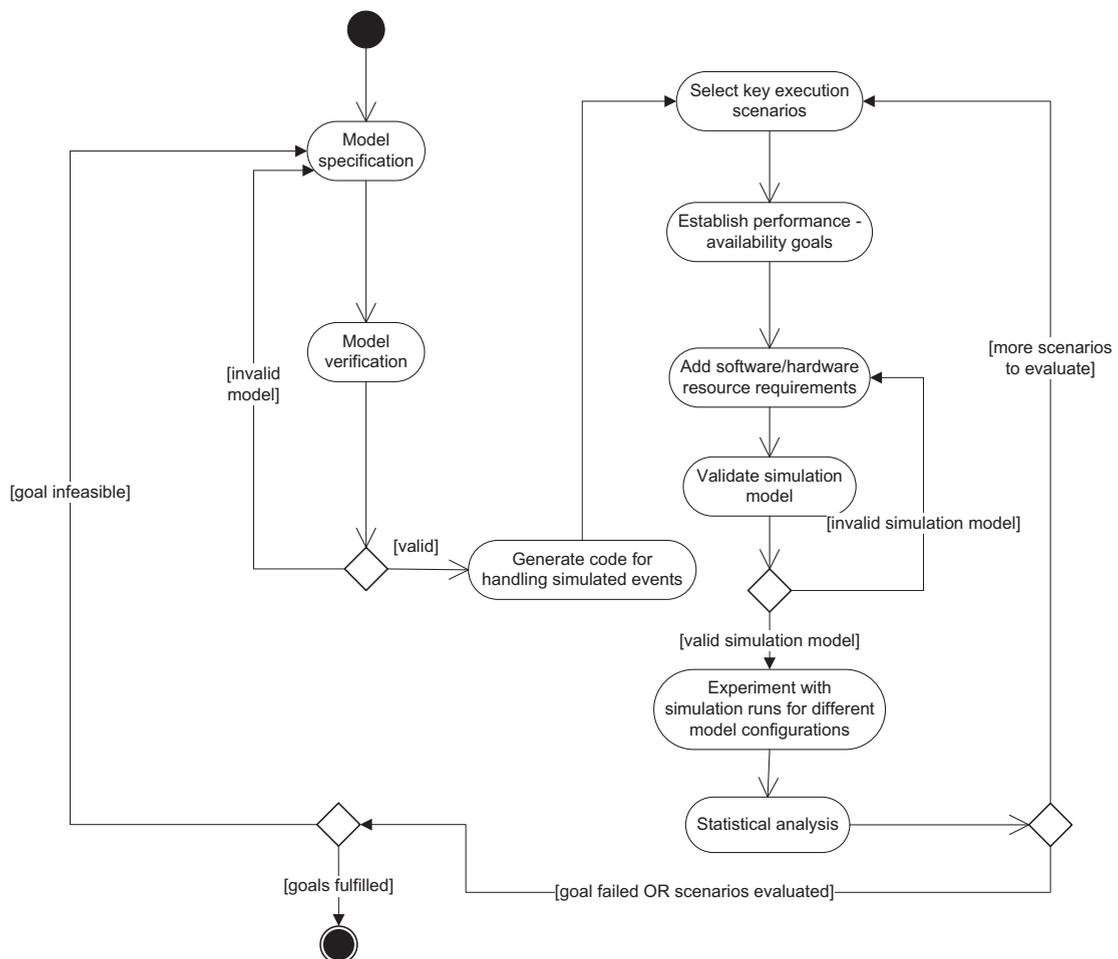

Fig. 1. Simulation-based evaluation of performance and availability goals for asynchronous system models.