

Simulation Modeling of UML Software Architectures

Simonetta Balsamo Moreno Marzolla

Dipartimento di Informatica

Università Ca' Foscari di Venezia

via Torino 155, 30172 Mestre (VE), Italy

e-mail: {[balsamo](mailto:balsamo@dsi.unive.it) | [marzolla](mailto:marzolla@dsi.unive.it)}@dsi.unive.it

February 21, 2003

Abstract

Quantitative analysis of software systems is being recognized as an important issue in the software development process. Performance analysis can help to address quantitative system analysis from the early stages of the software development life cycle, e.g., to compare design alternatives or to identify system bottlenecks. Modeling software systems by simulation allows the analyst to represent detailed characteristics of the system. We consider simulation for performance evaluation of software architectures specified by UML. We derive a simulation model for annotated UML software architectures. First we propose the annotation for some UML diagrams to describe performance parameters. Then we derive the simulation model by automatically extracting information about Use Case and Activity Diagrams from the XMI descriptions of UML diagrams. This information is used to build a discrete-event simulation model, which is finally executed. Simulation results are inserted back into the original UML diagrams as tagged values to provide feedback at the software architectural design level.

Keywords Discrete-Event Simulation, Software Architectures, Unified Modeling Language.

1 Introduction

In recent years it has been recognized that the software development processes should be supported by a suitable mechanism for early assessment of software performance. Early identification of unsatisfactory performance of Software Architecture (SA) can greatly reduce the cost of design change [22, 23]. The reason is that correcting a design flaw is more expensive the later the change is applied during the software development process. This is particularly true if the waterfall software development model is employed, as any change during the development process requires to start back from the beginning. However, this is still a relevant issue whenever a different software development process is used.

Both quantitative and qualitative analysis can be performed at the software architectural design level. Qualitative analysis deals with functional properties of the software system such as for example deadlock-freedom or security. Qualitative analysis is carried out by measurement or by modeling the software system to derive quantitative figures of merit, such as, for example, the execution profile of the software, memory or network utilization. We focus on performance models of software systems at the SA level. These include Queuing Networks [13], Stochastic Petri Nets [16] and Stochastic Process Algebra [12].

In this paper we consider quantitative evaluation of the performance of SA at the design level by means of simulation models. We consider SA expressed in terms of Unified Modeling Language (UML) [17] diagrams. We propose to annotate the UML diagrams using a subset of annotations defined in the UML Profile for Schedulability, Performance and Time Specification [18] (referred as *UML performance profile*).

Simulation is a powerful modeling technique that allows general system models, that is simulation models can represent arbitrarily complex real-world situations, which can be too complex or even impossible to represent by analytical models. We define a simulation model of an UML software specification introducing an almost one-to-one correspondence between behaviors expressed in the UML model and the entities or processes in the simulation model. This correspondence between system and the model helps the feedback process to report simulation results back into the original SA.

There are a few previous works dealing with simulation of UML specifications. Arief and Speirs [2, 3, 4] developed an automatic tool for deriving simulation models from UML

Class and Sequence diagrams. Their approach consists in transforming the UML diagrams into a simulation model described as an XML document. This model can then be translated into different kinds of simulation programs, even written in different languages. In this way the performance model is decoupled from its actual implementation. De Miguel et al. [10] introduced UML extensions for the representation and automatic evaluation of temporal requirements and resource usage, particularly targeted at real-time systems. The extensions are expressed in term of stereotypes, tagged values and stereotyped constraints. These were introduced in a commercial UML CASE tool, which has been made able to generate OPNET simulation models starting from annotated UML diagrams.

Previous simulation-based performance modeling approaches for evaluation of UML SA were developed before the UML performance profile was defined. Thus, they introduced their special extensions to UML or introduced non standard annotations to express quantitative information useful for deriving the model. In this paper we describe SAPS (Software Architectures Performance Simulator), a performance evaluation tool which translated UML Use Case and Activity diagrams into a discrete-event process oriented simulation model. The UML diagrams are annotated according to a subset of the UML Performance Profile [18]. This additional information is used to define the simulation model, which is finally executed. Simulation results are inserted back into the original UML diagrams as tagged values to provide feedback at the software architectural design level.

This paper is organized as follows. In Section 2 we illustrate the proposed methodology for generating simulation models from UML SA. In Section 3 we describe SAPS, a tool we built to implement that methodology. In Section 4 we illustrate a simple case study, and conclusions and future works are discussed in Section 5.

2 Overview of the methodology

In order to assist the software developer during the design process, we illustrate in Fig. 1 a general framework for quantitative analysis of UML SA [5]. The starting point is a description of the SA. We consider a description as a set of UML diagrams annotated with quantitative information in order to derive a simulation-based performance model. The model is

obtained using a suitable Modeling Algorithm. The model is then implemented in a simulation program, which is eventually executed. Simulation results are a set of performance measures, that can be used to provide a feedback at the original SA design level. The feedback should pinpoint performance problems on the SA, and possibly provide suggestions to the software designer about how the problem can be solved. The modeling cycle can be iterated until a SA with satisfactory performances is developed.

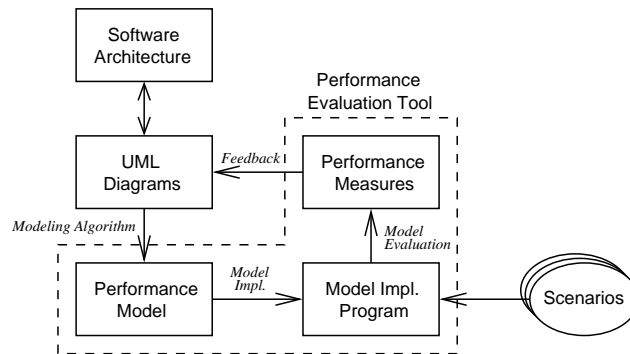


Figure 1: Framework for quantitative evaluation of UML diagrams

Note that the modeling and performance evaluation framework of Fig. 1 is independent from the particular performance model that we apply. In this paper we consider simulation-based performance models of UML SA. We describe SAPS, a prototype performance evaluation tool which processes an XMI [19] description of UML Use Case and Activity diagrams. The UML SA has to be annotated using a simplified subset of the UML Profile for Schedulability, Performance and Time Specification [18]. The simulation model is process oriented and its objects are derived by the analysis of the UML diagrams annotated with performance specification of the software system components. The simulation model is implemented as a discrete-event simulation program written in C++, whose execution provides results for a set of performance indices. We evaluate through simulation the mean response time associated with the execution of each scenario (Use Case) and each scenario step (Activity). Simulation results, i.e., the performance measures of the software components are inserted back into the original UML SA as tagged values to provide feedback to the system designer.

Figure 2 illustrates the structure of the performance simulation model derived from the UML diagrams.

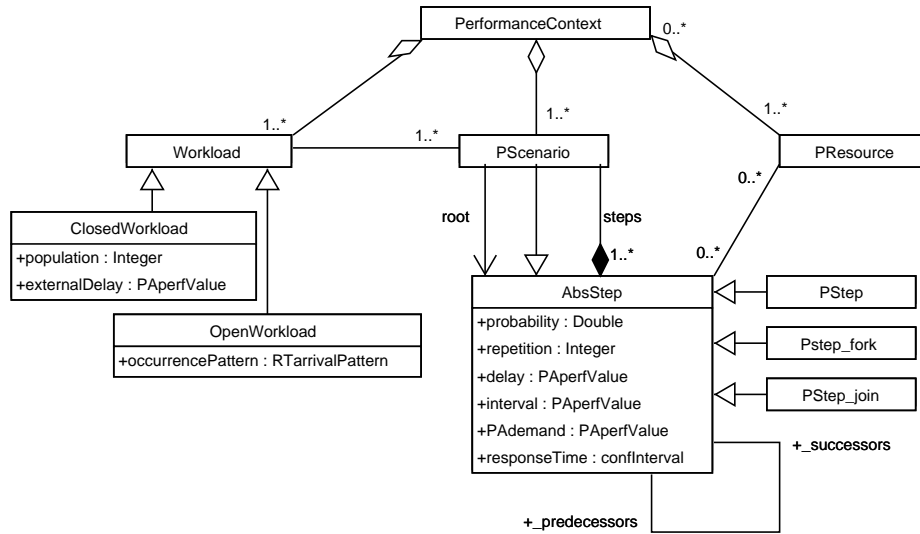


Figure 2: Structure of the simulation performance model

The basic object of the simulation model is a `PerformanceContext`. This object contains the other elements of the model, namely `Workloads` and `Scenarios`. `Workloads` can be open or closed, depending on whether the number of users accessing the system is unbounded or fixed. Closed workloads are characterized by two attributes (the exact notation used to describe the attributes is given in the next section):

population the total number of users in the workload

externalDelay (of type `PAPERfValue`, defined in Section 3) the delay between the end of a scenario execution and the beginning of the next request

Open workloads are characterized by the following attribute:

occurrencePattern (of type `RTarrivalPattern`, defined in Section 3) the pattern of inter-arrival times of consecutive requests

Each workload actually drives one or more scenarios. Each time a new user belonging to a workload requests service to the system, one of the scenarios associated with that workload is selected. Selection is done randomly, according to the probability associated to each scenario.

A scenario is a set of abstract scenario steps, represented by the `AbsStep` class. All

kinds of scenario steps are characterized by the following attributes:

probability the probability to execute this step, in the case the predecessor step has multiple successors

repetition the number of times this step has to be repeated

delay (of type `PAPerfValue`) an additional delay in the execution of this step, for example to model a user interaction

interval (of type `PAPerfValue`) the time between repetitions of this step, if it has to be repeated multiple times

PAdemand (of type `PAPerfValue`) the processing demand of this step

responseTime the computed delay between the starting and the finishing time of this step. The estimation of this quantity is the result of the simulation model execution. This value has type `confInterval`, which we define as the pair of the confidence interval bounds.

Abstract scenario steps can either be composite steps (described by the `PScenario` class), or atomic steps of different kinds. Scenarios are collections of steps; exactly one of these steps is marked as the root step (starting step) of the scenario. Atomic steps can be of type `PStep_fork` for nodes representing the creation of multiple execution threads, `PStep_join` for nodes representing synchronization points between different threads, and `PStep` for normal atomic steps.

There are two main differences between the performance model depicted in Figure 2 and the one described in the UML Performance Profile [18]. First, we assume a very simple model for resources: each Use Case has an associated computational resource; the resource is acquired with the Use Case starts, and is released at its completion. Thus, two instances of the same Use Case cannot be executed in parallel. This simplification is motivated by the fact that we evaluate SA at the architectural level, without assuming any implementation on a specific platform. Indeed we assume that the software architect ignores the specific hardware platform on which the software will be executed. Detailed resource modeling is usually done in later stages of the software development process.

Second, the structure we propose for the class hierarchy describing the processing steps

is slightly different from that in [18]. The UML Profile defines a `PScenario` class from which a `PStep` class is derived, thus every step is a scenario. This is because each step, at a deeper level of detail, could be modeled as a whole scenario, that is, a sequence of multiple sub-steps. We choose a different structure to model the `PStep` and `PScenario` hierarchy to keep atomic steps and scenarios as separate entities. We apply the Composite Pattern [11] to reflect the hierarchical nature of the processing steps. This choice makes the construction of the simulation model easier, because there are different kinds of step (e.g., `PStep`, `PStep_fork`, `PStep_join`) which are modeled as different simulation object types. The behavior of a fork step consists of activating all the successor steps concurrently. A join step waits for the completion of all predecessor steps before activating the successor. A normal step simulates execution according to the specified delays, and activates one of the successor steps. Finally, the behavior of a scenario is to activate its root step.

To summarize, the proposed approach SAPS to derive the simulation model to evaluate SA performance from UML Use Case and Activity diagrams is defined as follows:

1. Consider an UML representation of a software system in terms of Use Case and Activity diagrams. Both diagrams are respectively annotated as follows:
 - UML Use Case diagrams describe the interaction between the software system and one or more Actors requiring service. As proposed in [7, 20] we identify Actors to represent workloads applied to the system and Use Cases to represent scenarios. Actors can be stereotyped as `«PAopenLoad»` or `«PAClosedLoad»` to represent respectively open and closed population of users accessing the system. Use Cases are tagged with `PAprob` tags, whose value indicates the probability of executing that scenario.
 - Each Activity of an Activity diagram can be tagged with the following informations: the number of times the step has to be repeated (`PArep`); the delay between repetitions (`PAinterval`) of the same step; an additional delay for each step representing user “think time” (`PAdelay`); the service demand of the step (`PAdemand`).
2. The simulation model is automatically derived from the XMI description [19] of the UML diagrams. Currently SAPS uses the XMI dialect of the open-source ArgoUML

CASE tool [1], The simulation model is an instance of the general class structure depicted in Figure 2, and includes a `PerformanceContext` object and a set of workloads and scenarios.

3. The simulation model is executed, optionally asking the user to specify some parameters for the simulation, such as the desired confidence level for the estimation of the performance indices, the confidence interval width and the simulation length.
4. Simulation results are inserted back into the UML model as tagged values associated with Activities and Use Cases. We consider as simulation results the average delays (`PArespTime`) of Activities and Use Cases execution.

3 SAPS Tool Description

The steps of the proposed methodology are implemented into a simulation tool called SAPS. As concerns the first step of the annotation of UML diagrams, the UML Performance Profile suggests the use of TVL (Tag Value Language) to describe values of tags applied to model elements, which is a subset of the Perl language [24]. We use the freely available Perl interpreter library [8] to evaluate tag values, so taking advantage of the full Perl language. Note that both the UML Performance Profile and SAPS do not strictly depend on the specific language used to express annotations.

The `PAperfValue` and `RTarrivalPattern` data types are expressed according to the following BNF notation, which is a simplified version of the annotations defined in the UML Performance Profile:

```

< PAperfValue > := '['"assm"|"pred"|"msrd", "dist", < PDFstring >']'
< PDFstring > := '['< constantPDF > | < uniformPDF > |
< exponentialPDF > | < normalPDF >']'
< constantPDF > := < real >
< uniformPDF > := "uniform", < real >, < real >
< exponentialPDF > := "exponential", < real >, < real >
< normalPDF > := "normal", < real >, < real >

```


`< RTarrivalPattern > := '[' < bounded > | < unbounded > | < bursty >']'`
`< bounded > := "bounded", < integer >, < integer >`
`< bursty > := "bursty", < PDFstring >, < integer >`
`< unbounded > := "unbounded", < PDFstring >`

SAPS parses an XMI description of UML diagrams, annotated as described above. The UML model is translated into a C++ discrete-event simulation program according to step 2 of the proposed methodology. SAPS includes a general purpose discrete-event simulation library providing roughly the same functionality of the `Simulation` class of the SIMULA language [9], namely pseudo-parallel process execution using coroutines and Sequencing Set scheduling facilities. We developed the simulation library for several reasons, mainly code portability, availability of compilers and the necessity to use a freely available C library for parsing XML documents [15].

The simulation library includes random number generators and some statistical functions. Random variates of various distributions are generated using the uniform random number generator described in [14].

The library provides basic estimation functions, e.g. mean, variance and confidence interval. Since we consider steady-state simulation, we discard an appropriate initial portion of the observations to remove the initialization bias. The mean is computed using the method of independent replications [6]. The simulation stops when the relative width of the computed confidence intervals is smaller than a given threshold. Both the confidence level and the threshold can be defined by the user. If they are not provided, we assume default values of 90% confidence level and 5% threshold.

Each UML Actor is mapped into an appropriate `Workload` object, that can be an open or closed workload, depending on the stereotype which is applied to the UML element. Workloads are active objects in the simulation, which simply perform an endless loop in which they select and activate a Use Case, whose behavior is activating the root step of its associated Activity diagram. Each Activity in the Activity diagram is translated into the corresponding kind of `AbsStep` object. These objects simulate execution of the corresponding step or scenario according to the structure of the diagram and the value of the associated tags.

At the end of the simulation, the computed average response time of each `AbsStep` object (that is, a step or scenario) is inserted into the original UML diagram the `PArespTime` tag.


We propose to apply the SAPS tool as described in the modeling cycle illustrated in Figure 1. Namely, the software designer defines an UML SA with ArgoUML and specifies model parameters as tagged values associated with UML elements. Then, the user optionally selects parameters such as simulation length or confidence interval width, and runs the simulation. When the simulation finishes, the results are automatically inserted into the UML diagrams. At this point the software designer accesses the simulation performance results by opening the ArgoUML project file to access the results, and then possibly iterates the performance evaluation analysis by providing a new set of parameters in the UML model tags.

4 Case Study

In this section we illustrate with a simple example how SAPS works. The example involves an e-commerce application in which users can browse a web catalog of products or submit purchase orders. We assume an unlimited stream of users requiring service. Users interarrival time is exponentially distributed with mean A . Users can browse the online catalog with probability p , and make an order with probability $1 - p$. Browsing the catalog involves two sequential activities, which are: issuing a request to the product database and composing the web page. Making an order involves the following activities: selecting a product, filling the order form, processing the order and verifying the payment informations. Orders must be paid by credit card, whose number must be validated. Order validation and credit card checking are performed in parallel.


The UML Use Case and Activity diagrams are depicted in Figure 3. Model elements are tagged according to the notation described in Section 2.

Note that an analytical model of the system of Figure 3 can not be easily evaluated due to the fork/join component of the Make Order Activity. The simulation model which is derived from the SA is made of several active components (processes) arranged according to the



(a) Use Case

(b) Browse Catalog Activities



(c) Make Order Activities

Figure 3: UML representation of an E-commerce application.

structure of Figure 2. The actor is represented by an object of type `OpenWorkload`, which generates a stream of users. Use Cases are objects of type `PScenario` and have a queue which collects users according to their arrival order. Users select the Use Case to join according to the associated probability. Use Cases simulate one user request at a time activating the root step of the corresponding `PScenario`. When the scenario execution completes, a new user requests in the Use Case queue is started. Finally, there is one simulation process for each node in the Activity diagram. Nodes can be Activity nodes, which are modeled by

processes of type `PStep`, or fork and join nodes which are modeled by processes of type `PStep_fork` and `PStep_join`, respectively.

A numerical example of the performance results, computed as steady-state average delays are reported in Table 1. The results are obtained by setting $A = 50sec$ and $p = 0.9$; other model parameters are set as shown in Figure 3. The computed intervals are at 90% confidence level. The results are inserted into the original UML diagrams as values of the `PArespTime` tag. These tagged values are attached to each relevant model element. The software designer can now explore different situations by repeating the modeling and performance evaluation process with different tag values.

Scenario	Average delay (seconds)
Make Order	[2.98394, 3.00172]
Browse Catalog	[54.7787, 55.7254]

Table 1: Simulation results for the e-commerce application of Figure 3.

5 Conclusions

We have proposed a simulation-based performance modeling approach for UML software architectures. The notation we have used to describe performance parameters is a subset of the one defined in the UML Performance Profile [18]. We have derived a simulation model from Use Case and Activity diagrams. The simulation model is executed and the results are inserted into the original UML diagrams as tagged values. We have presented SAPS, a prototype tool to implement the methodology.

The proposed approach has been defined to evaluate software performances at the SA design level. We plan to extend this approach to further steps of the software development process, by considering Deployment diagrams and resource allocation. Further research will be devoted to a more complete set of performance measures. SAPS will be extended accordingly.

References

- [1] ArgoUML – Object-oriented design tool with cognitive support. <http://www.argouml.org/>.
- [2] L. B. Arief and N. A. Speirs. Automatic generation of distributed system simulations from UML. In *Proceedings of ESM '99, 13th European Simulation Multiconference*, pages 85–91, Warsaw, Poland, June 1999.
- [3] L. B. Arief and N. A. Speirs. Using SimML to bridge the transformation from UML to simulation. In *Proc. of One Day Workshop on Software Performance and Prediction extracted from Design*, Heriot-Watt University, Edinburgh, Scotland, November 25 1999.
- [4] L. B. Arief and N. A. Speirs. A UML tool for an automatic generation of simulation programs. In *Proceedings of WOSP 2000 [21]*, pages 71–76.
- [5] S. Balsamo, A. Di Marco, P. Inverardi, and M. Simeoni. Software performance: state of the art and perspectives. Technical Report MIUR SAHARA Project TR SAH/04, December 2002.
- [6] Jerry Banks, editor. *Handbook of Simulation*. Wiley–Interscience, 1998.
- [7] V. Cortellessa and R. Mirandola. PRIMA–UML: a performance validation incremental methodology on early UML diagrams. In *Proceedings of the Third International Workshop on Software and Performance (WOSP 2002)*, Rome, Italy, July 24–26 2002. ACM Press.
- [8] Comprehensive Perl Archive Network (CPAN). <http://www.cpan.org/>.
- [9] Ole-Johan Dahl and Kristen Nygaard. SIMULA—an ALGOL-based simulation language. *Communications of the ACM*, 9(9):671–678, September 1966.
- [10] M. De Miguel, T. Lambolais, M. Hannouz, S. Betgé-Brezetz, and S. Piekarec. UML extensions for the specifications and evaluation of latency constraints in architectural models. In *Proceedings of WOSP 2000 [21]*, pages 83–88.
- [11] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of reusable Object-Oriented programming*. Addison–Wesley, 1995.

- [12] Holger Hermanns, Ulrich Herzog, and Joost-Pieter Katoen. Process algebra for performance evaluation. *Theoretical Computer Science*, 274(1-2):43–87, 2002.
- [13] L. Kleinrock. *Queueing Systems*, volume 1. J. Wiley, 1975.
- [14] Pierre L’Ecuyer. Good parameters and implementations for combined multiple recursive random number generators. *Operations Research*, 47:159–164, 1999.
- [15] libxml: the XML C library for Gnome. Available at <http://xmlsoft.org/>.
- [16] M. Ajmone Marsan, G. Balbo, G. Conte, S. Donatelli, and G. Franceschinis. *Modeling with Generalized Stochastic Petri Nets*. J. Wiley, 1995.
- [17] Object Management Group (OMG). Unified modeling language (UML), version 1.4, September 2001.
- [18] Object Management Group (OMG). UML profile for schedulability, performance and time specification. Final Adopted Specification ptc/02-03-02, OMG, March 2002.
- [19] Object Management Group (OMG). XML Metadata Interchange (XMI) specification, version 1.2, January 2002.
- [20] R. J. Pooley and P. J. B. King. The Unified Modeling Language and performance engineering. In *IEE Proceedings – Software*, volume 146, pages 2–10, February 1999.
- [21] *ACM Proceedings of WOSP 2000, 2nd International Workshop on Software and Performance*, Ottawa, Canada, September 17–20 2000.
- [22] C. U. Smith. *Performance Engineering of Software Systems*. Addison-Wesley, 1990.
- [23] Connie U. Smith and Lloyd Williams. *Performance Solutions: A Practical Guide to Creating Responsive, Scalable Software*. Addison–Wesley, 2002.
- [24] Larry Wall, Tom Christiansen, and Jon Orwan. *Programming Perl*. O’Reilly & Associates, third edition, July 2000.