Integrated Functional and Performance Analyses of Concurrent Distributed Systems Described with the Language EMPA

M. Bernardo, L. Donatiello, R. Gorrieri

Dipartimento di Scienze dell'Informazione, Università di Bologna

The need of integrating the performance modeling and analysis of a concurrent system into the design process of the system itself has been widely recognized. The problem is that, in the case when quantitative aspects are neglected, time-critical concurrent systems (such as communication protocols) cannot be given completely satisfactory models and, moreover, these models cannot be used to estimate the system performance.

Unfortunately, it often happens that a concurrent system is first fully designed and tested for functionality, and afterwards tested for efficiency. The major drawback is that, whenever the performance is detected to be poor, the concurrent system has to be designed again, thereby negatively affecting both the design costs and the delivery at a fixed deadline. Another relevant drawback is that the tests for functionality and performance are carried out on two different models of the system, so one has to make sure that these two models are consistent.

In the last two decades a remarkable effort has been made in order to make existing formal description techniques suitable to support performance modeling and analysis by introducing the concept of time.

The most mature field where functional and performance aspects of concurrent systems are both considered is probably that of stochastic Petri nets, and in particular generalized stochastic Petri nets (GSPNs), since they have been extensively studied and successfully applied. They extend the expressiveness of classical Petri nets by associating with each net transition an exponentially distributed random variable specifying its duration.

The advantage of GSPNs is that functional and performance aspects are both taken into account since the beginning of the design process. These two aspects can then be separately analyzed on two different "projected models" (a classical Petri net and a Markov chain) obtained from the same "integrated model" (the GSPN), so we are guaranteed that the projected models are consistent. However two problems have to be addressed:

1. lack of compositionality, i.e. the capability of constructing nets by composing smaller ones,
2. inability to perform an integrated analysis, i.e. an analysis carried out directly on the integrated model without building projected models.

In order to solve the two problems above, we propose the adoption of a stochastic process algebra, since like classical process algebra, it naturally supplies compositionality. Such a stochastic process algebra extends the expressiveness of classical process algebras by representing each action as a pair composed of a type and a duration. Functional and performance analyses of a process term can be carried out on two consistent projected semantic models (a transition system labeled only on the type of the actions, and a Markov chain), as well as directly on the integrated semantic model (a transition system labeled on both the type and the duration of the actions) thanks to a suitable notion of “integrated equivalence”.

The stochastic process algebra is called Extended Markovian Process Algebra (EMPA). Its name stems from the fact that action durations are mainly expressed by means of exponentially distributed random variables (hence Markovian), but it is also possible to express actions having duration zero (hence Extended). The restriction to exponentially distributed durations simplifies the performance evaluation, as the performance model turns out to be a Markov chain. Also, such a restriction affects the semantic treatment, because the memoryless property of the exponential distribution allows us to define the integrated semantics for EMPA through the interleaving approach, in the same style as classical process algebras.