

Performance and Dependability Modeling with Stochastic Petri Nets

Organizer:

Heinz Beilner, Gianfranco Ciardo,
Christoph Lindemann, Kishor S. Trivedi

While measurement is a valuable option for assessing an existing system or a prototype, it is not a feasible option during the system design and implementation phases. Model-based evaluation has proven to be an attractive alternative in these cases. A model is an abstraction of a system that includes sufficient detail to facilitate an understanding of system behavior. Several modeling paradigms and various techniques for model evaluation are currently used in practice. The appropriate type of model depends upon the complexity of the system, the questions to be studied, the accuracy required, and the resources available for the study. Queueing network and simulation languages paradigms, in conjunction with corresponding analytic, numerical and simulative evaluation techniques have been employed for analyzing the performance of systems for more than 20 years. Due to recent developments in model generation and solution techniques and automated tools, large and realistic models can be developed and studied.

Recently, the use of stochastic Petri nets of various types has also been recognized as a useful modeling approach. The analysis of such Petri nets proceeds by utilizing the underlying continuous-time stochastic processes. Research work in stochastic Petri nets (SPNs) was begun about one decade ago by researchers from France, Italy, and the US. The first formal technical meeting constituted the first International Workshop on Petri Nets and Performance Models held in Torino, Italy in 1985. SPNs provide a unified modeling tool for both qualitative and quantitative analysis: In addition to available options for quantitatively evaluating corresponding Markov chains with appropriate numerical techniques, qualitative model properties can be studied employing a variety of well-established techniques for ordinary (untimed) Petri nets. Due to the availability of user-friendly software packages with graphical interfaces the development, modification, and quantitative evaluation of these SPNs is easier and less error-prone than, e.g., using a simulation language. These features constitute considerable advantages of SPNs over simulation languages and queueing networks.

In this seminar we concentrate on stochastic Petri nets, on the analysis of their underlying stochastic processes, and on their application to performance and dependability evaluation of computer systems, communication networks, and production systems. Particular topics to be discussed during the semi-

nar constitute efficient solution methods for the stochastic process underlying SPNs with deterministic and stochastic timing as well as hierarchical specification and solution techniques for SPNs. The aim of this seminar is bringing together the leading international researchers in this field with researchers from German academia and industry. The seminar is held in conjunction with the program committee meeting for the 6th International Workshop on Petri Nets and Performance Models.

Proceedings Editor: Martin Müller

Inhaltsverzeichnis

1	SPN Applications	6
2	Decomposition and Agregation Techniques	12
3	Related Modeling Formalisms	17
4	Dealing with General Distributions	21

1 SPN Applications

What the Designer of a Dependable System Expects of a Modeling Environment

Mario Dal Cin, Universit ät Erlangen–Nürnberg

The challenges of designing, e.g. embedded systems are: the care for high dependability and the joint development of hard- and software. These challenges require not only appropriate modeling support, they require also early feed-back from customers. To get this feed-back an early-virtual-prototype of the whole system is needed. Virtual prototyping entails co-modeling and co-evaluation. Co-modeling means that modeling starts from an abstract architecture of a joint hw/sw-system with a consistent semantics. GSPNs are among the best modeling techniques for dependable (fault-tolerant) systems. However, the integration of the existing (or future) GSPN modeling and evaluation tools with other tools is needed. Furthermore, the co-designer needs an easy access to GSPNs via familiar modeling paradigms and means to produce from his virtual prototype several orthogonal views – if GSPNs should become the modeling technique of a valuable co-design environment for dependable systems.

Modeling, Scheduling and Analysis of an Onboard Computing System

C. Girault, Universite P. et M. Curie (Paris 6)

L. M. Patnaik, Indian Institute of Science, Bangalore

We consider a two level decomposition of onboard realtime systems into repetitive major cycles, each one split in k minor cycles. The minor and major tasks are available at the beginning of their respective cycles and must be complete within them. All tasks have non deterministic execution times and are performed on a set of uniform processors. Minor tasks have priority on major ones which may only use the remaining time left in the k minor cycles.

A first colored stochastic net describes the dynamic management of minor tasks. Then it is embedded in second model for the scheduling and preemption of major tasks. The completion probabilities of the two sets of tasks are studied for several dynamic scheduling heuristics.

Petri Net Based Performance Engineering of Parallel Applications

Günter Haring, University of Vienna

This contribution presents the basic features of an approach (tool) supporting performance prediction of parallel systems (hardware and software), as it will be used in a performance oriented development cycle for parallel applications. The toolset is based on highly independent specifications of the workload (**p**rogram), the architecture (**r**esource) and the assignment of tasks/data out processing demands (**m**apping). The architecture of the toolset is based on a layer structure to generate and evaluate performance models, which have been generated automatically from the PRM-specification. The workload specifies the structured behaviour of the application (computation-communication structure) by acyclic task graphs. The behaviour of the architecture is represented by and is based on Petri Net building blocks. This approach allows an efficient hybrid simulation of the performance model, avoiding – or at least reducing essentially – the space problem, which is quite common on Petri Net models. In the two-level simulator a discrete event simulator, based on the task graphs triggers a special Petri Net simulator which instantiates and activates the basic building blocks in the appropriate way. The specification of building blocks at different levels of detail allow to trade off between modelling accuracy and evaluation speed. For a reasonable example a comparison between the performance prediction with the tool (PAPS) and actual performance measurements is given, the relative error of

the predicted execution time is less than 1%.

Modeling Real Time Systems with Stochastic Petri Nets

Günter Hommel, TU Berlin

Functional and quantitative requirements have to be taken into account for the construction of real-time systems. Essentially timeliness and dependability requirements can be distinguished. Where the first class of requirements leads mostly to deterministic models (time-out, sampling rate) dependability can only be captured using stochastic models. Additionally, in many soft real-time applications where quality of service is regarded, stochastic models are appropriate as well.

There is a trade-off between these two worlds of modeling. Using purely deterministic assumptions for a system allows to derive worst-case behaviour so that the timeliness can be shown. In this case the implicit assumption is that there are no unbounded events in the system. This means that all data rates and even the failure rates have to be bounded. This is, of course, not a valid assumption in the real world. Using stochastic assumptions on the other side prevents that the timeliness of the system can be guaranteed. Only probabilistic measures can be derived (timeliness is guaranteed to $9510e-10$).

For responsive systems that have to guarantee real-time behaviour and fault-tolerance a combination of both model worlds is essential. Deterministic and stochastic Petri nets (DSPNs) allow the modeling of such systems. A wide class of markovian and non-markovian nets has been intensively studied in our institute leading to the tool TimeNet. It supports steady state and transient analysis and simulation. For the simulation part several speedup techniques have been used (parallelization, variance reduction, RESTART). For a special class of applications in manufacturing a strongly restricted class of colored DSPNs has been defined that allows modular and hierarchical modeling. Due to the complexity problems it is still an open problem to handle real-time systems that show essentially deterministic behaviour with only few exceptions as e.g. the failure of the system. This field has to be studied in more detail in the future.

Modelling the Dependability of the French Air Traffic Control System Using GSPNs

Karama Kanoun, LAAS–CNRS Toulouse

The automated air traffic control system (CAUTRA) belongs to the category of systems that are real-time systems whose function are critical and demanding high level of dependability. The talk deals with the method that was followed to evaluate the availability of the system and to compose twelve alternative architectures.

The modelling approach is based on the derivation of two types of generic nets: constituent nets and dependency nets. The complete model of a given system architecture is gradually built up. Dependency nets represent for example: error propagation from a software component to an other software component, software stop following a hardware failure, sharing of a repair man by two hardware computers.

The advantages of this approaches are: the dependancy nets are generic and can be validated independly from each other (they are validated only with the constituent nets)

it is well adepted to model several architectures issued from the same specification differing by some fault tolerance strategies or maintenance falicity.

DSPN-Modelling of Usage Parameter Control in ATM-Networks

Bruno Müller-Clostermann, University of Essen

Traffic flow control mechanisms play an important role for the design and operation of future high-speed networks. Here we employ the class of Deterministic and Stochastic Petri Nets (DSPN) for the specification and evaluation of Usage Parameter Control (UPC) at the User Network Interface in ATM-networks.

After an overview of performance issues in traffic management DSPNs are applied for the specification of the functional and quantitative behaviour of traffic sources and control mechanisms. Traffic sources may be represented as stochastic Petri nets that model Interrupted Poisson Processes or the more

general class of Markov Modulated Poisson Processes in a straightforward way. Leaky bucket schemes (buffered and unbuffered) as well as window based techniques (jumping window, triggered jumping windows) can be specified precisely under inclusion of their time behaviour. Both the deterministic leaky rate and the fixed-length-sized window are modelled by deterministic transitions. One of the major objectives of UPC is to ensure a very low probability of cell losses for the traffic sources that are conformant with the parameters negotiated during the call set up phase. Numerical values for this so called violation probability have been computed by use of the tool DSPNexpress for different parameter settings.

DSPN have been shown to be a concise and unifying technique for the specification and analysis of Usage Parameter Control in ATM-networks. Due to the development of better techniques and tools in the area of timed and stochastic Petri nets future studies addressing the investigation of realistic scenarios should be feasible.

Buffer Sizing of ABR Traffic in an ATM Switch

Antonio Puliafito, Universita di Catania, Italy

M. B. Krishnan, K.S. Trivedi, Duke University, USA

I. Viniotis , NC State University, USA

The B-ISDN will carry a variety of traffic types: the Variable Bit Rate Traffic (VBR), Continuous bit rate traffic (CBR), Data Traffic and Available Bit Rate Traffic (ABR) that represents aggregate data traffic with very limited guaranties of quality. Of these VBR and CBR have timing constraints and need synchronous bandwidth; data traffic is relatively delay insensitiv. In this paper, we consider the VBR, Data and ABR traffic types and obtain the cumulative distribution function (cdf) of the queueing delay experienced by a burst in the output buffer of an ATM switch. The cdf is used to trade off buffer loss probabilities against deedline violation probabilities by adjusting the buffer size and delay deedline values. Large buffers result in low losses but queueing delays can become excessive and cause a high level of deedline violations. Both losses and violations are detrimental and an operating point must be chosen to achieve a balance. In this paper we study the nature of the

trade off. We develop a stochastic Petri net model assuming periodic burst arrivals for VBR and Poisson arrival processes for the Data and ABR traffic types at the burst level, and solve the model analytically (numerically) using a decomposition approach. This decomposition along with the inherent decomposability of the tagged customer approach for obtaining the cdf opens up the possibility of carrying out the computations for selecting the operating point each time that a call is admitted by using a parallel processor.

Modeling of a Distributed Multimedia System with Deterministic and Stochastic Petri Nets

Andreas Cramer, Essen

This contribution presents a case study for the performance analysis of the software implementation of a distributed multimedia system using a DSPN-type performance model. The goal of this model is the assurance of the realtime conditions of the system even in extreme situations and the quantitative performance analysis after model modifications that reflect system improvements. Given an existing implementation a functional model can be developed and the parameters that are needed to extend the functional model to a performance model can be measured. The developed functional Petri net model is a variation of the producer/consumer problem. This model is pretty large and unbounded so that analytical or numerical solution methods were not applicable. The transformation of the functional to a performance model can be divided into three tasks. First the parts of the functional model that deal with the “load” of the system have to be adapted. After that the processor scheduling policy “priority based preemptive resume” has to be modelled. Therefore model extensions were developed to avoid parallelism within and between processes. The third task is the introduction of measured delay distributions into the Petri net model. It turned out that there was no simple solution for this problem available, because of the preemptive resume processor scheduling policy. Facing that these extensions result in an explosion of the complexity of the model and that even the resulting huge performance model is not detailed enough to give precise answers to

our questions the extension from a functional to a performance model was not realised. In stead of the Petri net model a queing system model was developed that gives precise answers. Using this RESQ/ME model it could be shown that the end-to-end delay of the video transmission can be reduced from 770ms to 309ms without violating the realtime conditions of the system.

2 Decomposition and Agregation Techniques

Numerical Analysis of Hierarchical Stochastic Petri Nets

Peter Buchholz, Universit ät Dortmund

For many application areas Stochastic Petri Nets (SPNs) are an adequate formalism for quantitative modelling. Analysis of a SPN is usually performed by analysing the underlying Markov chain. Although Markov chain analysis is theoretically easy, most SPN models yield Markov chains of a size that cannot be handled even with todays' high performance computers. On possibility to deal with this complexity is to exploit some structural information for the generation and anlysis of a Markov chain underlying a SPN.

Hierarchical SPNs (HSPNs) are introduced as a means to describe systems in a convenient and highly structured way. HSPNs introduce a hierarchy by refining place-transition pairs. It can be shown that the resulting model structure is directly reflected in the structure of the state space and generator matrix of the underlying Markov chain. Exploitation of this structural information on Markov chain level yields several advantages. In particular, state spaces can be generated much more efficient and the size of models solvable with iterative numerical techniques is increased by an order of magnitude compared with conventional methods. Furthermore several aggregation and approximation techniques can be integrated naturally in the framework, allowing the, at least approximate, analysis of fairly large models.

A Decomposition Based Evaluation Method for Complex GSPN Models

Helena Szczerbicka, Universit ät Bremen

The major problem in the solution of GSPN models is a large state space of underlying Markov chain, which implies high memory and computational time requirements. It makes use of GSPN models difficult in many practical applications. A lot of methods focussing a reduction of computational complexity is available. However, most of them concentrate on a reduction of a complexity of an underlying Markov chain.

In our approach we follow a modeler point of view and propose a decomposition technique of GSPN models on the model level.

The model is decomposed and GSPN submodels are derived. The modeler composes then an overall model from submodels, taking into account interactions among subsystems. Fixed point iteration is required to cope with implications of cyclic situations when composing solutions of isolated submodels.

The important issue, how to decompose is solved by a computation of P-invariants of the overall model.

The basic concepts of composition are:

- a complementary transition reflecting the behaviour of the environment of the submodel and
- an interface for coupling submodels

The algorithm has been implemented and compared favorably concerning accuracy, convergence and a complexity reduction with an exact computation in several experiments.

Structured Modelling and State Space Reduction

Markus Siegle, Universit ät Erlangen–Nürnberg

In this presentation we look at different possibilities of how to approach

the state space problem when working with Markov models. We are particularly interested in modelling parallel and distributed systems with replicated components, i.e. scalable systems. An example for such a system is given using GSPN notation. We find that replicating subnets explicitly results in a state space growing exponentially with the degree of parallelism. In a dual representation where subnets are folded, we observe only linear growth. Next we look at examples for model transformation at the level of the high-level model description, using GSPN and stochastic process algebra (SPA) examples. We also address the problem of intelligent translation from the high-level formalism to the underlying CTMC. It is then shown how scalable systems can be described conveniently in a structured modelling framework where interacting submodels constitute the overall model. Submodels are typed, all submodels of one type being simply replicas of each other. The generator matrix for such a model is expressed with the help of tensor operations involving only small operand matrices. An efficient algorithm allows to compute matrices of a reduced stochastic process at the submodel type level, i.e. all submodels of a certain type are aggregated. Using this procedure, scalable models can be specified in a straight forward manner while one can still benefit from a reduced state space.

State Space Exploration of Superposed Generalized Stochastic Petri Nets (SGSPNs) Based on Structured Representations

Peter Kemper, Universität Dortmund

GSPNs are a modeling formalism which is well known for its suitability to describe concurrent systems; at least as well known is the state space explosion problem occurring in the analysis of GSPN models. In the context of performance analysis the synchronisation of formerly independent GSPNs via timed transitions into a then called SGSPN allows for a structured representation of the stochastic generator matrix of the underlying Markov chain. This relieves from the burden of the state space explosion problem but implies an additional overhead. The elimination of this overhead by a state space

exploration which is efficiently possible by exploiting the structured representation is extremely beneficial if not essentially necessary for a subsequent performance analysis.

Efficient Steady–State Solution of Markov Chains

Graham Horton, Universit ät Erlangen–Nuernberg

Timed Petri nets with exponential and phase–type firing distributions can be represented by a Markov chain. Steady state analysis requires the solution of a linear system of equations. Two characteristics of Markov chains leading to computational difficulties are NCDness and sidelength. We show why this is the case for standard iterative methods. We present a new method known as the “multi–level” algorithm, which is based on a stepwise, recursively applied aggregation of the Markov chain. We show why this method is susceptible to the two above mentioned difficulties and give some experimental results for the method applied to various markov chains.

A Structural Characterisation of Product Form Stochastic Petri Nets

Richard J. Boucherie, Universiteit van Amsterdam

Product form results for the equilibrium distribution of stochastic Petri nets are available in the literature. These results are based on assumptions for Markov chains describing the stochastic Petri net, and not on the structure of the Petri net. As the structure of the Petri net is one of the most important parts in the analysis of Petri nets, it seems natural to characterise the product form property on a structural level. We have provided part of

this characterisation: a necessary and sufficient condition on the structure of the network for a solution of the traffic equations to exist. The basis of this characterisation is the notion of minimal closed support T-invariants.

Probabilistic Evaluation of Large Markov Chains Using Uniformization and State Space Exploration

Boudewijn R. Haverkort, University of Twente
(Cooperation with Aad P. A. van Moorsel)

To evaluate transient measures for very large dependability models (in the Markovian context) uniformization is often called the method of choice. However, when the models of interest are very large memory problem might occur. We therefore propose to use only part of the state space of the Markov chain and compute measures from that. Of course these measures are not exact but they do provide bounds. In particular, by using two new variants of uniformization, i.e. orthogonal and partial uniformization, we are able to stepwisely increase the accuracy of our approach, i.e. we make the bounds tighter by stepwisely including more states in our computation. Depending on the situation at hand, this approach yields enormous memory gains. A question that remains is which states to include and which not. To answer this question we compared some exact results with state selection heuristics. It seems, from the case studies we performed so far, that an heuristic, in combination with partial uniformization yields very promising results.

A Petri Net Approach for the Performance Analysis of Business Processes

Alexander Schömig, Universität Würzburg

Recently, many companies have realized that their organization needs to be improved in order to enhance flexibility, efficiency, and effectiveness, i.e. the quality of services. Business Process Re-engineering has become the new organization paradigm to reconstitute profitability and competitiveness. Additionally, workflow management systems promise to provide a conceptual framework to describe and automate workflow. Analysts are asking for methods to evaluate alternative process designs for performance objectives, such as cycle times, throughputs and inventory levels. These measures are dependent on the dynamic behavior of a process and can hardly be derived by traditional methods, such as CPM or PERT.

We outline an approach to model the dynamic behavior of business processes by Stochastic Petri nets. We will summarize the problems we encountered following our approach to give directions for further research.

3 Related Modeling Formalisms

Stochastic Process Algebras and their Potential for the Integrated Design of Distributed Systems

Ulrich Herzog, Universität Erlangen–Nürnberg
(presented by Markus Siegle)

This talk is an introduction to the stochastic process algebra (SPA) formalism. We emphasize that constructivity is a major feature of process algebras. Constructivity means 1) composition of small descriptions in order to build more complex ones, 2) the possibility to abstract from the internal behaviour of a description, and 3) establishing equivalences between descriptions. The syntax of the SPA TIPP is explained, describing the meaning of the following operators: prefixing, choice, parallel composition, hiding and recursion. It is shown how SPA descriptions are translated into a labelled transition system (LTS) with the help of a set of deduction rules. Analysing such a semantic

model, functional, temporal and combined properties of the model can be derived. The notion of equivalence between descriptions is first defined at the level of LTS (bisimulation), and then lifted to the level of syntax. This lifting is called axiomatization, i.e. deriving a set of equational laws (this set must be both sound and complete) which allow to simplify a given SPA description at the syntax level. We conclude the talk by presenting the current state of tool support (PEPA-workbench [Hillsten and Gilmore, Edinburgh] and the TIPP-tool [Herzog et al., Erlangen])

Integrated Modeling Environment

Kishor Trivedi, Duke University, USA

On an attempt to improve the use of modeling techniques and tools in engineering practice, we propose the need for an integrated environment. In such an environment, an engineer can use an application specific interface to enter his modeling problem. The environment will include engines developed by different groups of developers and researchers based on possibly different modeling paradigms (fault tree, Markov chains, queueing networks, stochastic Petri nets etc.) Engines that could be integrated in the environment could be HIT, SHARP, SPNP, TomSpin The environment will include automatic detection of applicable engine and modeling paradigm and the automated translation into the appropriate engine's interface. The user will have access to a wide variety of engines and modeling paradigms without having to pay the overhead of learning details about them. A prototype of such an environment specialized to reliability modeling has been built by us for Boeing Commercial Airplane Company. The same idea could and should be carried through to function in a larger setting that includes more engines and many more application specific interfaces. An international project to support this idea is being proposed.

Decomposition solution of Large Networks of “Generalized Service Centers”

Murray Woodside, Carleton University, Canada

A scalable decomposition technique is given, suitable for very large GSPN systems, when the service center satisfy GSC conditions. They must be in a family generalized from state-machine by substitution of SISO subnet.

The condition make it easy to set up small auxiliary models one per center, of complexity slightly greater than the center itself. Computation of average performance measures was a fixed point iteration over the auxiliary models.

Approximation accuracies are of the order of 12 for examples tested.

Product Form Queueing Petri Nets

Falko Bause, Peter Buchholz, Universit ät Dortmund

The product form results for stochastic Petri nets are combined with the well known product form results for queueing networks in the model formalism of Queueing Petri nets yielding the class of Product form Queueing Petri Nets. This model class includes stochastic Petri nets with product form solution and BCMP Queueing Networks as special cases.

Exotic Algebras and Stochastic Petri Nets

Francois Bacelli, INRIA, France

I have shown that stochastic event graphs can be seen as $(\max,+)$ -linear systems in a random medium. A canonical representation was established, out of which two types of results were discussed:

- Taylor series expansions for the expected values of the canonical state variables in the case of open systems with Poisson input [joint work with V.Schmidt]
- parallel simulation algorithms based on the parallel prefix algorithm [joint work with M. Caroles]

The extension of this formalism to Free Chois Nets was briefly discussed.

4 Dealing with General Distributions

Markov Regenerative Stochastic Petri Nets with Age Memory Transitions

A. Bobbio, University of Brescia, Italy

We discuss a class of Markov Regenerative Stochastic PN (MRSPN) characterized by the fact that the stochastic process subordinated to the consecutive regeneration time points is a semi-Markov process with reward. This class of MRSPN can accommodate transitions with generally distributed firing times and associated memory policy of both enabling and age type. A unified analytical procedure is developed for the derivation of closed form expressions for the transient and steady-state probabilities.

Numeric analysis of Generalized Semi Markov Processes

Christoph Lindemann, GMD-FIRST

Numerical methods for discrete-event stochastic systems are needed in connection with performance and dependability models of computer and communication systems. We consider finite-state generalized semi-Markov processes with exponential and deterministic clock-setting distributions and provide an efficient numerical method for computing limiting distributions of such processes. The method is based on observation, at equidistant time points, of the continuous-state Markov process. The numerical method is valid not only when at most one deterministic event is active, but also when deterministic events may be concurrently active. The numerical technique is

applicable to networks of queues, deterministic and stochastic Petri nets, stochastic process algebras, and other discrete-event stochastic systems with an underlying process which can be represented as a generalized semi-Markov process.

This talk is based on joint work with Gerald Shedler and has been conducted while the speaker was a Visiting Scientist at the IBM Almaden Center in San Jose.

Analysis of Deterministic and Stochastic Petri Nets by the Method of Supplementary Variables

Reinhard German, Technische Universität Berlin

Stochastic Petri nets with both exponentially distributed and deterministic timing are well suited for the model-based performance and dependability evaluation. In this talk we present recent results on the numerical transient analysis of deterministic and stochastic Petri nets. The so-called “method of supplementary variables” is applied for the derivation of state equations which describe the temporal behavior. These equations consist of partial and ordinary differential equations combined with initial and boundary conditions.

Two different numerical techniques are presented for the solution of the state equations:

- an iterative algorithm based on discretization and
- an easier solution for a special case.

These techniques have been implemented and added to the software tool TimeNET. TimeNET is a tool which supports the design and evaluation of stochastic Petri nets with and has been developed at the Technische Universität Berlin. TimeNET is especially designed for dealing with non-Markovian stochastic Petri nets. A tool demonstration will be given during the talk in order to illustrate the new transient analysis component.

Transient Distributions of Cumulative Reward

Edmundo de Souza e Silva, University of Rio de Janeiro,
Brazil

Markov reward models have been used to solve a wide variety of problems. In these models, reward rates are associated to the states of a continuous time Markov chain and impulse rewards are associated to the transitions of the chain. Several methods have been developed in the past to solve reward models. These include techniques based on Laplace transform, partial differential equations and uniformization. We briefly review a few of the methods used to calculate transient distributions of cumulative rewards. We then present a new efficient algorithm to calculate this measure when both rate and impulse rewards are present in the model. The development is based only on probabilistic arguments and the recursions obtained are simple and have a low computational cost.

Importance Sampling in UltraSAN

Bill Sanders, University of Illinois, USA

Model-based evaluation of reliable systems is difficult due to the complexity of these systems and the nature of the dependability measures of interest. The complexity creates problems for analytical model solution techniques and rare events make traditional simulation methods inefficient. Importance sampling is a well-known technique for improving the efficiency of rare event simulations. However, finding an importance sampling strategy that works well in general is very difficult. This fact motivated the development of an environment for importance sampling that supports a wide variety of model characteristics and measures. The environment is based on stochastic activity networks, and importance sampling strategies are specified using the new concept of the “governor”. The governor supports dynamic importance sampling strategies by allowing the stochastic elements of the model to be refined based on the evolution of the simulation. Several examples of the technique were presented.