

From *Proceedings of the Third Int'l. Workshop on Petri Nets and Performance Models*, Kyoto, Japan, December 11-13, 1989, pp. 74-84.

## REDUCED BASE MODEL CONSTRUCTION METHODS FOR STOCHASTIC ACTIVITY NETWORKS\*

W. H. Sanders<sup>†</sup> and J. F. Meyer<sup>‡</sup>

<sup>†</sup>The University of Arizona  
Computer Engineering Research Laboratory  
Department of Electrical and Computer Engineering  
Tucson, AZ 85721

and

<sup>‡</sup>The University of Michigan  
Computing Research Laboratory  
Department of Electrical Engineering and Computer Science  
Ann Arbor, MI 48109

### ABSTRACT

During the past ten years, significant progress has been made in the development and application of stochastic extensions to Petri nets. A principal focus of this work has been models that can support system evaluation with respect to various measures of performance, dependability, and performability. Several such model classes now exist and, as evidenced by their use in specific evaluation studies, they are capable of representing the kind of complex behavior exhibited by contemporary distributed computer architectures and computer-communication networks. However, a number of problems remain open, particularly those associated with the evaluation of large-scale systems. Specifically, the issue here is the size and complexity of the stochastic process, derived from the underlying net model, which serves as a "base model" for subsequent solution of the measures in question. If this base model constructed by standard means, e.g., it is identified with the marking behavior of the net, traditional methods of solution quickly become intractable for large systems, limiting their application to systems of only moderate complexity. This problem is addressed in the context of stochastic activity networks (SANs) by developing base model construction methods that account for symmetries in SAN structure and are tailored to the variable (e.g., response time, time to failure, etc.) in question. We find that such a technique can yield dramatic reductions in state-space size while preserving stochastic properties required for practical means of solution. Moreover, unlike state aggregation methods that rely on explicit knowledge the more detailed state space, this technique permits direct construction of a reduced base model, thus avoiding size limitations associated with more traditional approaches to model simplification. The presentation includes the both the theory required to establish the technique's validity and some illustrations of its effectiveness in the evaluation of a CSMA/CD computer network.

---

This work was supported in part by the Office of Naval Research under Contract No. N00014-85-K-0531.

## I INTRODUCTION

Large-scale distributed systems pose challenging evaluation problems due to their user requirements, complexity, and scale. In particular, demands for both high performance and high dependability are common. This, coupled with the fact that resources that are distributed, results in systems which exploit parallel processing. Furthermore, due to a balance of demands for high performance and dependability, such requirements typically call for systems which exhibit degradable performance in the presence of faults, i.e., between extremes of nondegraded performance (as would be exhibited if the system were fault-free throughout utilization) and fully degraded performance (system failure), the system is able to perform at intermediate levels which provide some benefit to the user.

These characteristics, namely parallelism, fault tolerance, and degradable performance, must be accounted for in the evaluation process, and hence place significant constraints on the types of models that may be used. Specifically, while queueing networks have been used extensively for the evaluation of computing systems [1, 2], their application to systems that exhibit complex parallelism (e.g., contention for multiple resources), fault tolerance, and degradable performance is not straightforward. Furthermore, queueing networks cannot easily represent fault-related behavior, since the probabilistic nature of their structures (e.g., branching probabilities) is fixed. On the other hand, stochastic Petri nets (SPNs) [3, 4] and subsequent extensions thereof are much better suited to the modeling of systems that exhibit such properties. In particular, the use of timed “transitions” or “activities” with probabilistic timing permits, via different interpretations of tokens, simultaneous representation of characteristics related to both performance and dependability.

Because of this capability, over the past ten years, considerable effort has been devoted to the development of SPN-type models for the evaluation of complex distributed systems (see, for example, the work documented in [5, 6]). More precisely, if the property being evaluated (e.g., throughput rate, response time, time to failure, etc.) is denoted by a random variable, say  $Y$ , an SPN-type model determines a stochastic process  $X$  which, in turn, is the basis for determining the probabilistic nature of  $Y$ . Accordingly,  $X$ , in this context, is referred to as a *base model* (that supports the solution of  $Y$ ; see [7], for example). This research has been particularly fruitful and has resulted in the definition of formal model classes that are well suited to generating base models for systems of the type described above. As a consequence, process models that, in the past, were either hand-constructed Markov models or were determined by complex hand-written simulation code, could now be generated by

simple, understandable stochastic network models. However, a number of problems remain open, particularly those associated with the evaluation of large-scale systems. When SPN-derived base models are constructed by standard means, traditional methods of solution quickly become intractable for large systems, limiting their application to systems of only moderate complexity.

In particular, when the solution methods to be employed are analytic, construction of an associated process model (base model) typically assumes that the state space is the set of all possible reachable, stable (or tangible) markings of the net. Likewise, when simulation methods are used to solve the model, each of the transitions or activities is taken to be different event type in the associated simulation program. In the case of large-scale systems, the resulting base models can become intractably large in terms of state-space size, precluding analytic solutions, or the number of event types, precluding accurate simulations. Efforts aimed at solving this problem have been limited. In particular, prior to the work described below, neither the structure of the net nor the nature of the performance variable has been accounted for in attempts to construct simpler base models.

If stochastic extensions to Petri nets are to become a truly useful basis for system evaluation, there is a need for techniques that permit systems (represented in this manner) to be evaluated without intermediate generation of a detailed stochastic process or simulation program. This need was recognized by the research community in a panel session entitled “Applications of Performance Petri Net Models: Where does the future lie?” at the *Second International Workshop on Petri Nets and Performance Models*. In effect, what is needed are developments, analogous to product-form solutions for queueing networks, that permit evaluation of a large system to be based on special knowledge of the network model and performance variable. This knowledge can then be used to reduce the amount of information required to obtain the desired evaluation results or, in short, to “solve” the probabilistic nature of the performance variable.

To date, work toward this goal has proceeded in several directions. In particular, one approach has been to look for subclasses of nets that have the “product-form” property. While some results have been obtained [8, 9], they appear to be limited to very restricted subclasses of nets. Another approach has been through the use of colored tokens (e.g., [10]) or high-level Petri nets [11]. A third approach, and the one we have taken, makes use of both the structure of the stochastic Petri net and the desired performance variables to generate a “smaller” stochastic process that is both solvable and yields the desired information

regarding the performance variable. The approach is based on the use of stochastic process lumping theorems, but makes use of the specified performance variable and structure of the network to select an appropriate notion of state and does not require generation of the complete marking space of the network. Furthermore, the generalized notion of state employed permits the determination of “activity-related” as well as “marking-related” behaviors. This permits one to ask questions regarding the number of activities that complete during an interval.

Preliminary work with this goal was documented in [12], under the name “variable driven construction methods”. In [12], we considered a single class of variables and a subclass of stochastic activity networks which consisted of one or more subsystems, each of which could contain, in turn, one or more replications of an even smaller subsystem. This work generalizes that of [12] in several significant ways. First, we consider many different variable types, capturing transient as well as steady-state system behaviors. Second, we consider a more general multi-level stochastic activity network structure, wherein the construction operations can be applied any number of times to build arbitrarily complex stochastic networks. Finally, we present an algorithm to construct the “reduced” stochastic process without requiring the generation of the complete marking space.

The remainder of this paper is organized as follows. In the following section, we briefly review the basic definitions concerning stochastic networks and then define the types of variables to be employed in the construction process. These variables, as mentioned earlier, can be used to estimate both transient and steady-state system characteristics. The third section describes the construction operations used, contains theorems stating the validity of the method, and presents a procedure for generating the “reduced” base model stochastic process for a given stochastic activity network and performance variable. This is followed by some examples which both illustrate the method and demonstrate its effectiveness in reducing the size of a state space.

## II NETWORK-LEVEL MODELS AND VARIABLES

### A Stochastic Activity Networks

*Stochastic activity networks* (SANs) [13, 14] incorporate features of both stochastic Petri nets and queueing models. Structurally, SANs have primitives consisting of *activities*, *places*, *input gates*, and *output gates*. Activities (“transitions” in Petri net terminology) are of two types, *timed* and *instantaneous*. Timed activities represent activities of the modeled

system whose durations impact the system’s ability to perform. Instantaneous activities, on the other hand, represent system activities which, relative to the performance variable in question, complete in a negligible amount of time. Cases associated with activities permit the realization of two types of spatial uncertainty. Uncertainty about which activities are enabled in a certain state is realized by cases associated with intervening instantaneous activities. Uncertainty about the next state assumed upon completion of a timed activity is realized by cases associated with that activity. Places are as in Petri nets. The use of gates permits greater flexibility in defining enabling and completion rules.

The stochastic nature of the nets is realized by associating an *activity time distribution function* with each timed activity and a *probability distribution* with each set of cases. Generally, both distributions can depend on the global marking of the network. A *reactivation function* [14] is also associated with each timed activity. This function specifies, for each marking, a set of *reactivation markings*. Informally, given that an activity is activated in a specific marking, the activity is *reactivated* whenever any marking in the set of reactivation markings is reached. This provides a mechanism for restarting activities that have been activated, either with the same or different distribution. This decision is made on a per activity basis (based on the reactivation function), and is not a net-wide execution policy.

The execution of stochastic activity networks is discussed detail in several places, including [15]. Informally, though, SANs execute in time through completions of activities that result in changes in markings. More specifically, an activity is chosen to *complete* in the current marking based on the relative priority among activities (instantaneous activities have priority over timed activities) and the activity time distributions of *enabled* activities. A case of an activity chosen to complete is then selected based on the probability distribution for that set of cases. These two choices determine uniquely the next marking of the network, which is then obtained by executing the input gates connected to the input of the activity chosen and the output gates connected to the chosen case. This procedure is repeated by considering the activities enabled in the new marking.

In order to solve a stochastic activity network by analysis for a particular variable, a stochastic process must be constructed that both supports the variable and is solvable. Traditionally, for both SANs and other variants of stochastic Petri nets, this has been accomplished by taking the possible stable (tangible) markings as states (see for example, [15] where this is called the “marking behavior”). This approach results in a large state space that can support variables which can be written in terms of possible markings of the

net. A second process, known as the “activity-marking behavior”, was considered in [15]. In this case, the states of the process (known as *am-states* in that which follows) are taken to be pairs of the form  $(a, \mu)$ , where  $a$  is a timed activity and  $\mu$  is a next stable marking which may be reached upon completion of  $a$ . More precisely then, the *activity-marking behavior* (or *am-behavior*) of a SAN is a stochastic process  $(R, T, L) = \{R_n, T_n, L\}$  where  $R_n$  is the state reached after the  $n$ th timed activity completion,  $T_n$  is the time of the  $n$ th activity completion and  $L$  is the total number of the transitions of the process. This process, although typically larger than the marking behavior for a given SAN, can capture activity as well as marking behaviors.

While the activity-marking behavior can conceptually support an extremely large class of performance variables, solution of the probability of occurrence of arbitrary sets of trajectories is often intractable. Accordingly, we opt for solutions based on higher-level (less detailed) models which support more restricted classes of variables. To make these distinctions more precise, a stochastic process which supports a large class of variables (e.g., activity-marking behavior and marking behavior) is referred to as a *detailed base model*; a stochastic process which is constructed specifically to support a designated performance variable is a *reduced base model*. The following section presents several classes of variables which can be used in the construction of reduced base models.

## B Performance Variable Specification

As outlined in the introduction, our method makes use of knowledge regarding the desired performance variables as well as the structure of the net to aid in the construction process. In order to make this possible, it is necessary to formally specify the “types” of performance variables that may be considered. Previous work regarding “reward models” [16] (and associated “reward variables”) provides an instructive example in this regard. Informally, a reward model consists of a stochastic process and a “reward structure”. The reward structure relates possible behaviors of the process to a specified performance variable. Typically, this is done by associating a “reward rate” with each state, its interpretation being the rate at which reward accumulates while the process is in that state. In this case, the performance variable is taken to be the reward accumulated over some utilization interval (either finite or infinite). By associating different reward rates to states, one can construct performance variables with many different interpretations.

The network-level structure of a stochastic activity network (or other stochastic Petri

net extension) allows us to define the reward structure at the network level, rather than the state level. This approach has several distinct advantages over the state-level approach outlined above. First, the assignment of rewards and interpretation of solutions are more natural, because it is done at the level at which the modeler thinks. Second, since rewards are assigned at the network level, they can be used in the construction procedure. Further advantages of this approach, along with a careful development of the underlying theory, can be found in [17] where we construct a general reward structure at the SAN level and systematically generate variables from that reward structure. For the purposes of this paper, however, it suffices to briefly list the variables defined so that they can be used in the construction procedures developed in the next section.

In particular, the variables considered are organized according to reward structure type, category within a reward structure type, and variable type within a category. At the highest level, the variables are distinguished by the choice of reward structure type. By type we mean one or more classes of functions that have a particular interpretation in terms of the SAN. For a given reward structure type, variables can be further distinguished by the interval of time that they depend on. Three categories of variables are distinguished at this level. The first category, *instant-of-time variables*, represents the status of the SAN at either a particular time  $t$  or in steady state. The second category, the *interval-of-time variables*, represent particular aspects of the behavior of a SAN over some finite or infinite interval of time. The final category, the *time-averaged interval-of-time variables*, represent particular aspects of the behavior of a SAN over some interval divided by the length of the interval.

Three types of variables are considered for the interval-of-time and time-averaged interval-of-time variable categories. The first type represents the total or time-averaged reward (relative to a particular reward structure) accumulated during some interval  $[t, t + l]$ . The second type corresponds to an interval of length  $l$  as  $t$  goes to infinity, and is useful in representing the reward that is accumulated during some interval of finite length in steady-state. The final variable type corresponds to the total or time-averaged reward accumulated during an interval starting at  $t$  and of length  $l$  as  $l \rightarrow \infty$ .

A SAN-level reward structure, together with eight variable types, has been defined that quantifies benefits associated with activity completions and particular numbers of tokens in places. In particular, we define an “activity-marking oriented reward structure” as follows:

**Definition 1** *An activity-marking oriented reward structure of a stochastic activity net-*

work, with places  $P$  and activities  $A$ , is a pair of functions:

$\mathcal{C} : A \rightarrow \mathbb{R}$  where for  $a \in A$ ,  $\mathcal{C}(a)$  is the reward obtained due to completion of activity  $a$ , and

$\mathcal{R} : \mathcal{P}(P, \mathbb{N}) \rightarrow \mathbb{R}$  where for  $\nu \in \mathcal{P}(P, \mathbb{N})$ ,  $\mathcal{R}(\nu)$  is the rate of reward obtained when, for each  $(p, n) \in \nu$ , there are  $n$  tokens in place  $p$ ,

where  $\mathbb{N}$  is the set of natural numbers and  $\mathcal{P}(P, \mathbb{N})$  is the set of all partial functions between  $P$  and  $\mathbb{N}$ .

Informally, impulse rewards are associated with activity completions (via  $\mathcal{C}$ ) and rates of reward are associated with numbers of tokens in sets of places (via  $\mathcal{R}$ ). An element  $\nu \in \mathcal{P}(P, \mathbb{N})$  is referred to as a *partial marking*. The marking is partial in the sense that natural numbers are assigned to some subset of  $P$ , namely the domain of the partial function  $\nu$ . This assignment is made in a manner identical to the way a (total) marking assigns natural numbers to all the places in the set  $P$ . Although  $\mathcal{R}$  has a countably infinite domain, the number of elements  $\nu$  that are of interest to the modeler and, hence, deserving of a non-zero reward assignment will generally be small compared to, say, the number of reachable stable markings of the SAN. Similarly, it will usually be the case that only a fraction of the SAN's activities will have non-zero rewards associated with their completions. We thus adopt the convention that rewards associated with activity completions and partial markings are taken to be zero if not explicitly assigned otherwise.

Variables are then defined by writing functions in terms of the reward structure and behavior of the network, and fall into the three categories discussed above. In particular, two instant-of-time variables are defined. The first quantifies the behavior of a stochastic activity network at a particular time  $t$ . More precisely, if we let  $V_t$  denote this variable type then

$$V_t = \sum_{\nu \in \mathcal{P}(P, \mathbb{N})} \mathcal{R}(\nu) \cdot I_t^\nu + \sum_{a \in A} \mathcal{C}(a) \cdot I_t^a,$$

where

$I_t^\nu$  is an indicator random variable representing the event that the SAN is in a marking such that, for each  $(p, n) \in \nu$ , there are  $n$  tokens in  $p$  at time  $t$ , and

$I_t^a$  is an indicator random variable representing the event that activity  $a$  is the activity that completed most recently at time  $t$ .



When  $I_t^\nu$  and  $I_t^a$  converge in distribution for all  $\nu$  and  $a$  with non-zero rewards as  $t$  approaches  $\infty$ , the “steady-state” reward obtained at an instant of time can be studied. If we denote the random variable with this steady-state distribution as  $V_{t \rightarrow \infty}$ , its value can be expressed as

$$V_{t \rightarrow \infty} = \sum_{\nu \in \mathcal{P}(P, \mathbb{N})} \mathcal{R}(\nu) \cdot I_{t \rightarrow \infty}^\nu + \sum_{a \in A} \mathcal{C}(a) \cdot I_{t \rightarrow \infty}^a,$$

where

$I_{t \rightarrow \infty}^\nu$  is an indicator random variable representing the event that the SAN is in a marking such that, for each  $(p, n) \in \nu$ , there are  $n$  tokens in  $p$  in steady-state, and

$I_{t \rightarrow \infty}^a$  is an indicator random variable representing the event that activity  $a$  is the activity that completed most recently in steady-state.

As mentioned earlier, three variable types are considered for the interval and time-averaged-interval variables corresponding to an interval of length  $l$  starting at time  $t$  ( $[t, t + l]$ ), an interval of length  $l$  as  $t \rightarrow \infty$  ( $[t, t + l], t \rightarrow \infty$ ), and an interval starting at  $t$  as  $l \rightarrow \infty$  ( $[t, t + l], l \rightarrow \infty$ ). Variable types of the interval category are denoted by “ $Y$ ” while variables types of the time-averaged category are denoted by “ $W$ ”, each with the appropriate subscript. In particular, let

$$Y_{[t, t+l]} = \sum_{\nu \in \mathcal{P}(P, \mathbb{N})} \mathcal{R}(\nu) \cdot J_{[t, t+l]}^\nu + \sum_{a \in A} \mathcal{C}(a) \cdot N_{[t, t+l]}^a, \text{ and}$$

$$W_{[t, t+l]} = \frac{Y_{[t, t+l]}}{l}$$

where

$J_{[t, t+l]}^\nu$  is a random variable representing the total time that the SAN is in a marking such that, for each  $(p, n) \in \nu$ , there are  $n$  tokens in  $p$  during  $[t, t + l]$ , and

$N_{[t, t+l]}^a$  is a random variable representing the number of completions of activity  $a$  during  $[t, t + l]$ .

Similarly, if the required variables converge in distribution, we can define

$$Y_{[t, t+l], t \rightarrow \infty} = \sum_{\nu \in \mathcal{P}(P, \mathbb{N})} \mathcal{R}(\nu) \cdot J_{[t, t+l], t \rightarrow \infty}^\nu + \sum_{a \in A} \mathcal{C}(a) \cdot N_{[t, t+l], t \rightarrow \infty}^a,$$

and

$$W_{[t,t+l],t \rightarrow \infty} = \frac{Y_{[t,t+l],t \rightarrow \infty}}{l},$$

where

$J_{[t,t+l],t \rightarrow \infty}^\nu$  is a random variable representing the total time that the SAN is in a marking such that, for each  $(p, n) \in \nu$ , there are  $n$  tokens in  $p$  during a interval of length  $l$  in steady-state, and

$N_{[t,t+l],t \rightarrow \infty}^a$  is a random variable representing the number of completions of activity  $a$  during an interval of length  $l$  in steady-state, and

$$Y_{[t,t+l],l \rightarrow \infty} = \sum_{\nu \in \mathcal{P}(P, \mathbb{N})} \mathcal{R}(\nu) \cdot J_{[t,t+l],l \rightarrow \infty}^\nu + \sum_{a \in A} \mathcal{C}(a) \cdot N_{[t,t+l],l \rightarrow \infty}^a,$$

and

$$W_{[t,t+l],l \rightarrow \infty} = \lim_{l \rightarrow \infty} \frac{Y_{[t,t+l]}}{l}$$

where

$J_{[t,t+l],l \rightarrow \infty}^\nu$  is a random variable representing the total time that the SAN is in a marking such that, for each  $(p, n) \in \nu$ , there are  $n$  tokens in  $p$  during  $[t, \infty)$ , and

$N_{[t,t+l],l \rightarrow \infty}^a$  is a random variable representing the number of completions of activity  $a$  during  $[t, \infty)$ .

As shown in [17], these variable types can be specialized, through appropriate choices of reward structures, to cover a wide variety of specific performance, dependability, and performability measures.

### III REDUCED BASE MODEL CONSTRUCTION

Given a performance variable of the type described above, let us now consider the problem of constructing a SAN-derived stochastic process that i) supports solution of the variable in question and ii) is solvable by practical means. As stated in the introduction, for stochastic extensions of Petri nets, the state space of such a process is typically taken to be the to the set of reachable, stable markings of the net, i.e., the resulting process (base model)

is the net’s *marking behavior*. Certain variables, however, cannot be supported by this process, e.g., variables referring to activity completions (or, in SPN terms, transition firings) which, in general, cannot be uniquely inferred from knowledge of a marking trajectory. Accordingly, a wider variety of performance variables can be supported if, along with the marking, the state includes the name of the most recently completed activity; in this case, the resulting base model is referred to as the *activity-marking behavior*. While the latter choice suffices to support the types of variables described in the previous section, their use with realistic systems can result in state spaces that are very large. For certain classes of systems, however, reduced base models can be constructed that are much smaller (in number of states) and still preserve system information necessary to solve for the chosen performance variables. This section discusses the theory and application of reduced base model construction methods. These methods are useful for hierarchical systems where, at certain levels in the hierarchy, subsystems are replicated and where contributions to total reward (as defined by an activity-marking variable) are the same for identical subsystems. Such hierarchical systems are becoming increasingly more prevalent and, in a computer context, include both multiprocessor systems (with many identical processors and busses) and computer networks (with many identical stations).

As with the conditions on system structure, the conditions on the performance variables are not severe. In fact, since the structure of replicated systems is identical, it is natural to assume that possible behaviors of identical subsystems would be treated in an equivalent manner. An example of this would be a computer network consisting of many types of stations, where one is interested in the queue length at a particular type of station, not at an individual station within a type. The next subsection introduces operations that aid in the construction of stochastic activity networks that meet these conditions.

## A Construction Operations

In order to insure that models built have characteristics that permit construction of a reduced base model, we construct them in a bottom-up manner using two operations. Before we can do this, however, a formal definition of the operands of the construction operations is needed. Specifically, we introduce the notion of a “SAN-based reward model”.

**Definition 2** *A SAN  $S$  with places  $P$ , together with a reward structure  $(\mathcal{C}, \mathcal{R})$  and set of distinguished places  $P_D \subseteq P$ , is a SAN-based reward model  $(S, \mathcal{C}, \mathcal{R}, P_D)$ .*

A SAN-based reward model (SBRM) provides us with a structure that can be used to construct larger models that have compact state-space representations. Since the operations discussed below are defined on SAN-based reward models, they operate on the reward structure as well as the network. This insures that the defined variable is supported by the reduced model.

The first construction operation *replicates* a SAN-based reward model. This operation is useful when a system has two or more identical subsystems. The effect of the operation depends on a set of places which is a subset of the distinguished places of the input SBRM. These places allow communication between the replicated submodels and are not replicated when the operation is carried out. Informally, the result of the replicate operation on a SAN-based reward model is another SAN-based reward model, where

1. The SAN component is constructed by replicating the original SAN a certain number of times, holding some subset of the distinguished places of the input SBRM common to all replicate SANs, and
2. The reward structure is constructed by assigning an impulse reward to each replicate activity equal to its reward in the original model and reward rates to each partial marking in the new model equal to the rate assigned to the corresponding partial marking in the original model, and
3. The set of distinguished places is the set of places used in the construction operation.

The replicate operation allows us to construct SAN-based reward models that consist of several identical component SBRMs. This permits the representation of systems that consist of a single replicated structure. Typically, however, distributed systems consist of several different structures, each of which may be replicated. The combination of several different structures is accomplished using the join operation. As with the replicate operation, the *join* operation both acts on and produces SAN-based reward models. Informally, the effect of the operation is to produce a new SAN-based reward model which is a combination of the individual subnetworks and reward structures. Again, certain places play an important role in the construction operation. In this case, however, a *list* of places is associated with each component SBRM in a manner such that the cardinality of each list of places is the same.

In the joined model, the corresponding places in each list form a single place, with connections to gates and activities as in the individual models. These lists of places thus

allow communication between joined subnetworks. As with the replicate operation, each list of places must be a subset of the set of distinguished places of the component SBRM. The reward structure for the new SBRM is constructed in a similar manner to that used for the replication operation. Specifically, the reward structure is constructed by 1) assigning a reward to each activity in the new model equal to the reward of the corresponding activity in the original model and 2) assigning a reward rate for each partial marking in the new model equal to the rate assigned to the corresponding partial marking in the original model.

These operations can be applied iteratively to build models of systems that are organized in a hierarchical manner. We refer to a SAN-based reward model constructed in this way as a *composed SAN-based reward model*. In order to describe the structure of a composed SAN-based reward model, we make use of a directed tree with three types of nodes: *model/reward structure nodes*, *replicate nodes*, and *join nodes*. Model/reward structure nodes have a degree of zero. These nodes define distinct subnetworks and reward structures in the composed model, and serve as a basis on which to apply the construction operations. Replicate nodes have a degree of one with an outgoing arc that points to the SAN-based reward model that is replicated. Join nodes have a degree equal to the number of SBRMs that are joined together. In this case, each arc points to a SBRM that is joined.

Information contained inside each node describes the parameters of the particular operation. In particular, each model/reward structure node contains a triple,  $(S, \mathcal{C}, \mathcal{R})$ , specifying the subnetwork and reward structure for the node. The SBRM associated with the node consists of the SAN and reward structure of the node together with a set distinguished places equal to the set of places of the SAN. Replicate nodes contain an integer denoting the number of times the associated SBRM is to be replicated (given as  $N$  in the figure) and a set of places to be used in the replicate operation. Join nodes contain an integer (denoting the number of SBRMs to be joined) and a list of places for each SBRM to be joined.

## B Representations of State

We now consider possible state space and, in turn, base model, representations for composed SAN-based reward models. One approach would be to use activity-marking or marking states as the state representation, and a detailed base model as the process representation. This, however, leads to excessively large state spaces for many models. The aim, then, is to find one or more alternative representations for state that lead to reduced base models. If support of the variable in question were the only concern, the set

of possible values of the variable could be taken as the states of the process. However, the state behavior of this process is not generally Markovian or discrete state, and thus very difficult to solve. On the other hand, if “solvability” were the only issue, the states of the detailed base model could be lumped into a single state in a manner such that the behavior of the reduced base model was easy to deduce. This will not suffice, however, since this process will not generally support a selected variable. Thus both support and solvability are important considerations in selecting a reduced base model representation.

Since we consider a process to be solvable if it is Markov, results regarding lumpings of a Markov process [18] and, equivalently, functions of a Markov process that are Markov [19] are relevant. These results give conditions on a Markov process and a lumping of that process (or equivalently, a functional of that process) such that the new process is also Markov. These conditions are limited, however, in two respects. First, they consider only the solvability of the resulting process; no requirement is made that the resulting process support any particular variable. Second, a direct application would require that the detailed process be generated and that a suitable lumping be found. Machine memory size limitations typical preclude generation of the detailed process and, even if the process could be generated, there is no easy way (without additional information regarding the system being modeled) to determine a suitable lumping from the detailed model.

Stochastic activity networks and the construction operations introduced in the previous subsection provide us with a way to avoid both of these difficulties, due to the way in which the construction operations were defined. In particular, the order of application of the operations (as specified by the directed tree representation introduced in the previous subsection) determines a notion of state which both supports the variable in question and is Markov whenever the detailed base model representation is Markov. Informally, this notion of process state is determined such that for each replicate operation, the number of replicate SBRMs in each possible submodel “state” is recorded and for each join operation, a vector of the “state” of each joined submodel is kept. At the lowest level, the “state” of a SBRM model is the marking of the submodel. The complete state for the composed model is then taken to be a pair where the first component is the impulse reward of an activity in the model and the second component is a state formed as described above.

To give a more precise example of how the notion of state would be formulated for a particular application of the replicate and join operations, consider the composed SAN-based reward model of Figure 1, where the lowest level nodes are arbitrary models and

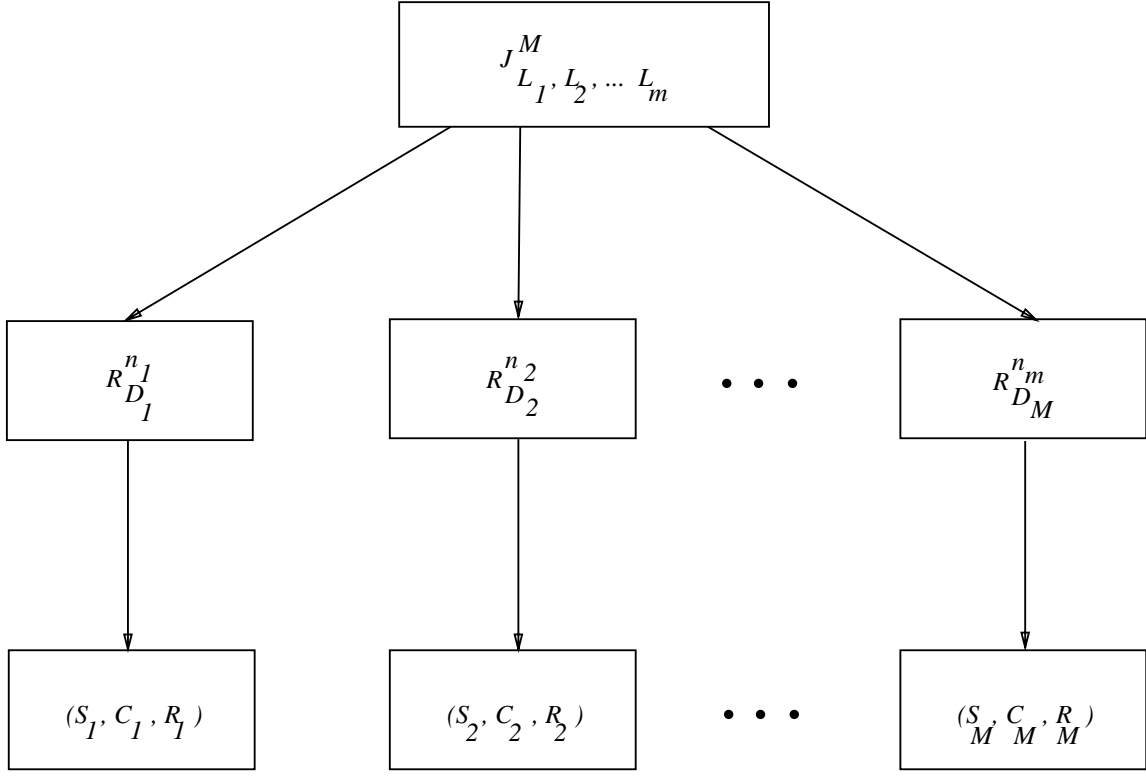


Figure 1: Example Composed SAN-Based Reward Model

reward structures. Now let  $E$  be the state space of the activity-marking behavior for this model, and  $M$  to be the set of marking components of this am-space. Then define projections to examine the markings of submodel types and individual submodels. Specifically, for each  $\mu \in M$  define  $\mu^j$  to be the projection of (global marking)  $\mu$  on the places of submodels of type  $j$ ,  $j = 1$  to  $m$ . Similarly, define  $\mu^{j,k}$  to be the projection of (global marking)  $\mu$  on the places of submodel  $k$  of type  $j$  (where places within a submodel have the same ordering).

Furthermore, for each  $\mu^j$  let  $B_{\mu^j} = \langle \mu^{j,1}, \mu^{j,2}, \dots, \mu^{j,n_j} \rangle$  be the bag of projected markings for submodels of type  $j$  in marking  $\mu$ . The bag formalism [20] is useful here, since it allows us to compactly characterize the different configurations various replicates of a submodel are in, without distinguishing an ordering among them. As will be seen in the following, this representation allows us to specify a functional which, due to the nature of the replicate and join operations, can serve as a reduced base model for the system in question. If, further, we let

$$f(a, \mu) = (\mathcal{C}(a), B_{\mu^1}, B_{\mu^2}, \dots, B_{\mu^m}) \quad (1)$$

then this mapping defines a functional of the am-behavior associated with the SBRM of Figure 1. More specifically, this functional is defined in terms of two processes derived from the am-behavior: the minimal process and the discrete-time imbedded process. Following the convention used for Markov renewal processes, the minimal behavior  $Z = \{Z_t \mid t \in \mathbb{R}^+\}$  of the am-behavior is a process defined such that  $Z_t$  is the am-state of the network at time  $t$ . Likewise, the discrete-time imbedded process associated with the am-behavior is a process  $R = \{R_n \mid n \in \mathbb{N}\}$  defined such that  $R_n$  is the am-state reached after the  $n$ th timed activity completion. Given these processes, the minimal and discrete-time imbedded behaviors of a reduced base model can be defined as

$$U = \{f(Z_t) \mid t \in \mathbb{R}^+\}$$

and

$$T = \{f(R_n) \mid n \in \mathbb{N}\}$$

respectively.

While the functional notion of Equation (1) is convenient for two levels of operations, it becomes cumbersome if extended directly to more levels. Instead, we use a set of equations relating SAN-based reward models at one level to those at the next lower level. The nature of this relationship depends on the specific operation used at the particular level. As with the two-level example just given, each replicate node corresponds to a “number of” relation in the state space, and is represented as a bag of submodels at the next lower-level. On the other hand, the “state” of each joined SBRM must be represented explicitly, and hence the operation is represented as a vector of joined submodels. At the lowest level are the submarkings of the individual SAN-based reward models. As with the two-level case discussed above, activities are named in a way which designates the particular submodel that they belong to, and projections of markings on particular submodels are designated by superscripts on the marking (i.e.,  $\mu^{n_1, n_2, \dots, n_h}$  denotes the projection of global marking  $\mu$  on the places of the  $n_1$ th SBRM of the highest-level operation, the  $n_2$ th SBRM of the next-level operation, and so on).

For example, consider the SAN-based reward model of Figure 2. The functional describing the mapping from am-states to reduced states is written as:

$$\begin{aligned} f(a, \mu) &= (C(a), V) \\ V &= (B_1, \mu^2) \end{aligned}$$



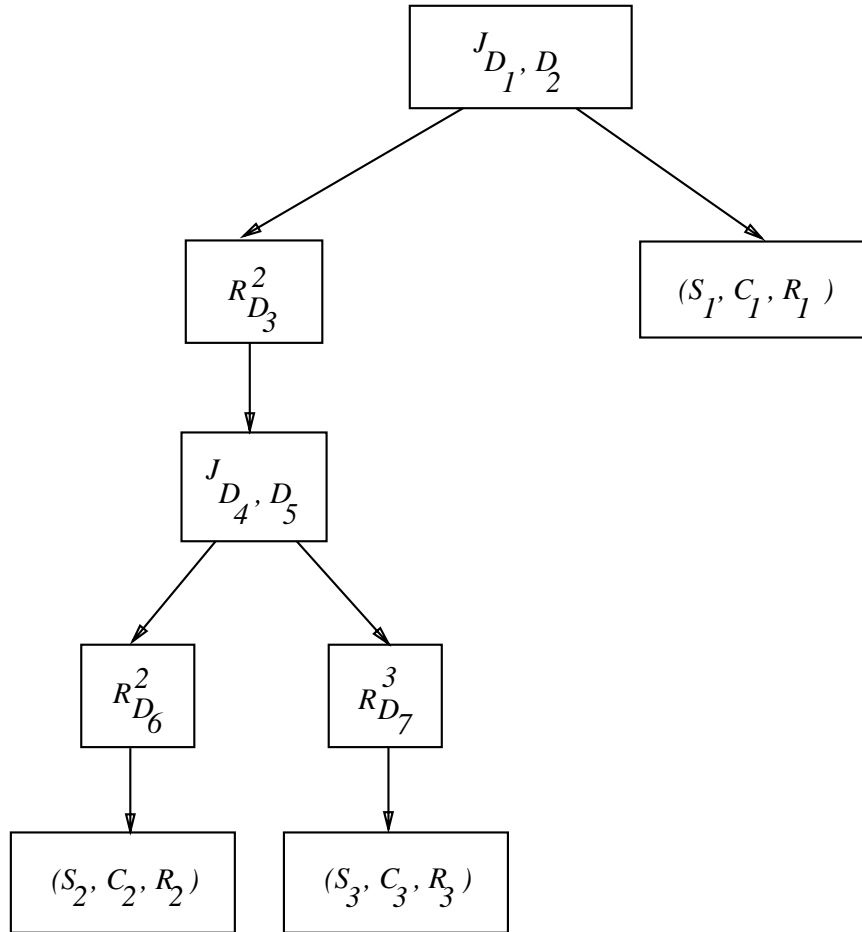


Figure 2: Example Composed SAN-based Reward Model

$$\begin{aligned}
B_1 &= \langle V_{11}, V_{12} \rangle \\
V_{11} &= (B_{111}, B_{112}) \\
V_{12} &= (B_{121}, B_{122}) \\
B_{111} &= \langle \mu^{1111}, \mu^{1112} \rangle \\
B_{112} &= \langle \mu^{1121}, \mu^{1122}, \mu^{1123} \rangle \\
B_{121} &= \langle \mu^{1211}, \mu^{1212} \rangle \\
B_{122} &= \langle \mu^{1221}, \mu^{1212}, \mu^{1223} \rangle .
\end{aligned}$$

Using this notation, the letter  $B$  denotes a bag of SBRMs at the next lower level and  $V$  denotes a vector. The superscripts on markings denote the marking of a particular submodel.

As alluded to earlier, in order for reduced base models to be useful, they must support the variables of interest and be solvable. Sufficient conditions for this to occur are known, and are now given. In particular,

1. The reduced base model of a composed SBRM supports the instant-of-time, interval-of-time, and time-averaged interval-of-time activity-marking oriented variables defined herein,
2. The minimal behavior ( $U$ ) of the reduced base model of a composed SBRM is Markov whenever its minimal am-behavior is Markov, and
3. The discrete-time imbedded behavior ( $T$ ) of the reduced base model of a composed SBRM is Markov whenever its am-behavior is Markov.

Theorems and accompanying proofs that establish these results can be found in [15], along with theorems and proofs that show when the detailed base models are Markov. Intuitively, however, these models are solvable because replicate submodels of a particular type behave in an equivalent manner when in identical markings. Similarly, since the performance variable was constructed in a bottom up manner as part of the composed SAN-based reward model, it is “symmetric” with respect to different replicates of a given submodel type and thus supported by the reduced base model. Since we can generate these processes without first generating a detailed base model, we have avoided both of the problems associated with classical lumping methods.

Taken together, these theorems show that both the minimal and discrete-time imbedded behaviors of the reduced base model of a composed SBRM are Markov whenever the corresponding activity-marking behavior is Markov. Since these processes also support the variable of interest, they can be used to solve for the variable’s probabilistic behavior. Construction procedures which generate a reduced base model are now investigated.

## C Construction Procedure

While a reduced base model is indeed solvable whenever the corresponding detailed base model is and, moreover, it supports the specified performance variable, these facts do not suggest a manner to generate a reduced base model. In fact, the proof of solvability given in [15] relies on lumping arguments. Clearly, one would hope that a procedure could be formulated that did not rely on the prior generation of a detailed base model. This is, in fact, possible, and the resulting procedure is more efficient than those derived for detailed base models.

The procedure operates on reduced states and “representative” markings. A *representative marking of a reduced state* is any marking that corresponds (via the functional defined earlier) to the marking portion of the reduced state. Once a state has been expanded, the representative marking is no longer needed; hence representative markings need be associated with states only until they are expanded.

In the following procedure,  $f$  is the functional defined earlier which maps am-states to reduced states,  $U$  is the set of reduced states,  $X$  is the set of reduced states yet to be expanded and representative markings for these states,  $T$  is the set of transition rates from marking portions of reduced states to reduced states. Furthermore,  $Y$  is a set of triples denoting possible next reduced states, representative markings for these states and rates to these states from the marking portion of a reduced state. The marking portion of a reduced state  $u$  is denoted as  $u^m$ . The rate from a marking portion of a reduced state,  $u^m$ , to a reduced state  $u'$  is written as  $\lambda_{u^m, u'}$ .

**Procedure 1** (*Procedure to generate the reduced base model of a stochastic activity network with initial marking  $\mu_0$ .*)

Let  $U = \{f(\nabla, \mu_0)\}$ .  
 Let  $X = \{(f(\nabla, \mu_0), \mu_0)\}$ .  
 Let  $T$  equal the null set.  
 While  $X \neq \text{null}$ :

Let  $(u, \mu) = \text{some\_element}\{X\}$ .  
 Let  $X = X - \{(u, \mu)\}$ .  
 Let  $Y = \text{poss\_}U(\mu)$ .  
 For each  $(y_u, y_\mu, y_\lambda) \in Y$ :  
   If  $y_u \in U$  then  
     Let  $T = T \cup \{(u^m, y_u)\}$ .  
     Let  $\lambda_{u^m, y_u} = y_\lambda$ .  
   else  
     Let  $U = U \cup \{y_u\}$ .  
     Let  $T = T \cup \{(u^m, y_u)\}$ .  
     Let  $\lambda_{u^m, y_u} = y_\lambda$ .  
     If there does not exist a  $u \in U$  such that  $u^m = y_u^m$   
       Let  $X = X \cup \{(y_u, y_\mu)\}$ .  
 Next  $(y_u, y_\mu, y_\lambda) \in Y$ .  
 While end.

Algorithm 1 (*poss\_U*) generates the set of next possible reduced states and rates to these states from a particular marking portion of a reduced state. In this algorithm, the activity that is completed is referred to as the *representative activity*. The number of activities that are in the set of replicate activities is denoted by  $n$ . After possible next stable markings and probabilities of these states for each representative activity are computed using an algorithm given in [15], rates to these states are then found by multiplying the rate due to a single activity by the number of replicate activities in the particular submodel marking. The following algorithm describes this in a more precise manner.

**Algorithm 1** (*Generates the set,  $Y$ , of reduced states, representative markings, and rates to these states from a representative marking  $\mu$  of a particular reduced state.*)

Let  $Y$  equal the null set.

Let  $D$  be the set of pairs  $(a, n)$  where  $a$  is a representative activity and  $n$  is cardinality of the set of replicate activities of  $a$  in  $\mu$ .

For each  $(a, n) \in D$ :

Compute the set of possible next stable markings reached upon completion of  $a$  in  $\mu$  (i.e.  $NS(a, \mu)$ ), the probability of reaching each of these markings (i.e.  $h_{\mu, a}$ ), and check whether this distribution is invariant over possible sets of activity choices  
 If the distribution is not invariant over possible sets of activity choices then

Signal *SAN* is not well specified and abort algorithm.

For each  $\mu' \in NS(a, \mu)$ :

If  $\exists (y_u, y_\mu, y_\lambda) \in Y$  such that  $f(a, \mu') = y_u$  then  
 $(y_u, y_\mu, y_\lambda) = (y_u, y_\mu, y_\lambda + r_a(\mu) * n * h_{\mu, a}(\mu'))$ .

else

Let  $Y = Y \cup \{(f(a, \mu'), \mu', r_a(\mu) * n * h_{\mu, a}(\mu'))\}$ .

*Next*  $\mu' \in NS(a, \mu)$ .  
*Next*  $(a, n) \in D$ .

Note that in order for this algorithm to run successfully to completion, the SAN in question must be “well-specified”. Informally, a SAN is well specified if its probabilistic behavior is completely specified, i.e. an unambiguous probabilistic representation can be derived from the SAN. This may not be the case if more than one instantaneous activity is enabled in a reachable marking. An algorithm to check whether a specific SAN is well specified is given in [15].

An example illustrating the effectiveness of reduced base model construction methods is now considered.

#### IV EXAMPLE CSMA/CD LOCAL AREA NETWORK

The protocol employed is a variant of non-persistent CSMA/CD. When a station has something to send, it waits an exponentially distributed amount of time before attempting the transmission. After this delay, it senses the channel. If the station detects an idle channel, it begins to transmit. One of two events will then occur: either 1) the channel was actually idle, so the transmission begins normally, or 2) a collision occurs, since another station had begun transmitting, but its signal had not propagated to the first station at the time the channel was sensed. If a collision occurs, it is cleared after an exponentially distributed amount of time. If the station detects a busy channel, it returns to the wait state and after the delay period once again attempts to gain access to the channel.

Two scenarios are considered. In the first, all stations in the network are identical. In the second, one station is considered to be a high-priority station, and all other stations are considered to be normal priority. Here, the priority of a station is determined by the mean time it waits before trying to access the bus; the high-priority station waits (on the average) one-tenth the time of the other stations.

A stochastic activity network representing a single station connected to the channel is given in Figure 3. The gate and activity parameter values for the model are given in Tables 1 and 2, respectively. The station represents either a high or normal priority station, depending on the duration of the expected wait between attempts to access the channel. Figure 4 gives a SAN representing the network submodel. Since the goal here is to illustrate the state space savings that can be achieved, a detailed description of the correspondence

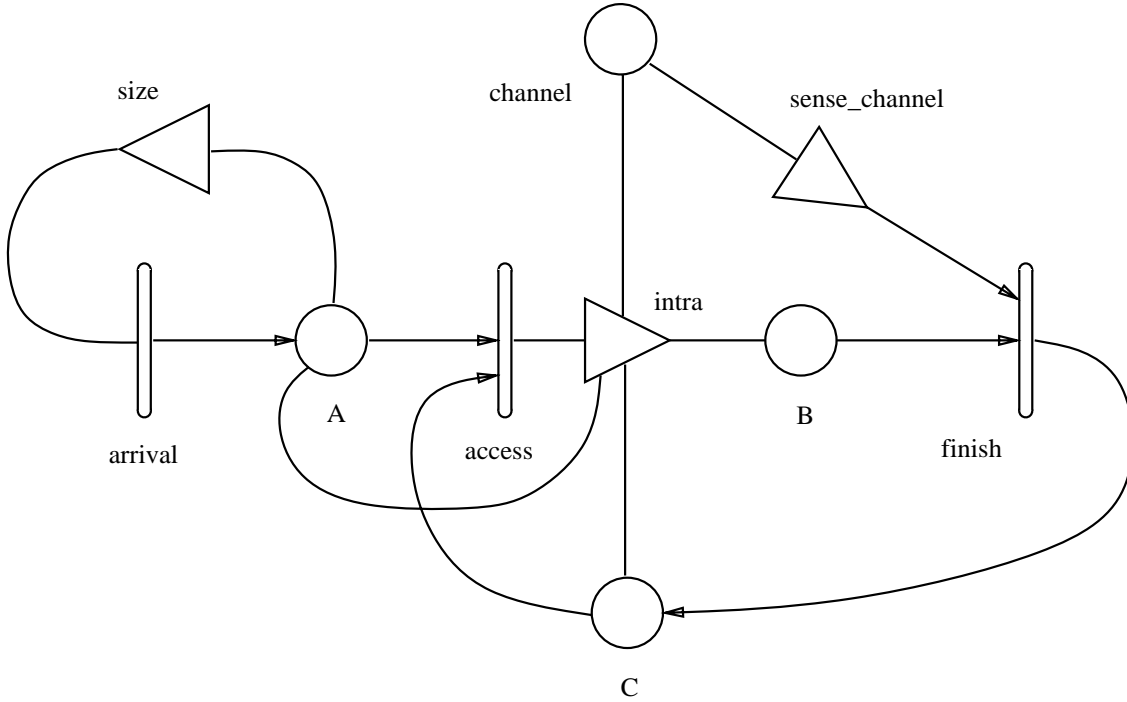


Figure 3: Station Submodel

between the submodels and protocol is not given.

Reward structures for the various performance variables are now considered. The first performance variable studied is the expected queue length in steady-state at each type of station. This can be obtained using the reward structure

$$\mathcal{C}_{st}(a) = 0 \quad \forall a \in A$$

$$\mathcal{R}_{st}(\nu) = \begin{cases} i & \text{if } \nu = \{(A, i)\} \\ 0 & \text{otherwise} \end{cases}$$

and variable  $E[V_{t \rightarrow \infty}]$ . For the homogeneous system, the expected queue length at an individual station is then  $\frac{E[V_{t \rightarrow \infty}]}{N}$ , where  $N$  is the number of stations of the given type. Since markings and activity completions in the network submodel do not contribute to this variable, the reward structure for this submodel ( $\mathcal{C}_Z, \mathcal{R}_Z$ ) is defined such that  $\mathcal{C}_Z(a) = 0, \forall a \in A$ , and  $\mathcal{R}_Z(\nu) = 0, \forall \nu$ .

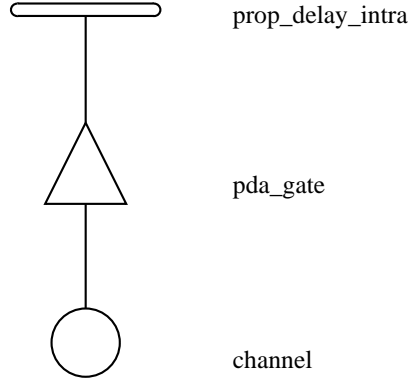
Performance variables representing the fractions of time various actions are present on the channel are constructed using a null reward structure for each station submodel (i.e,

Gate	Type	Enabling Predicate	Function
<i>size</i>	input	$MARK(A) < 2$	identity
<i>intra</i>	output	–	if ( $MARK(channel) == 0$ ) { $MARK(channel) = 1$ ; $MARK(B) = 1$ ; } else if ( $MARK(channel) == 1$ ) { $MARK(channel) = 3$ ; $MARK(C) = 1$ ; $MARK(A) = MARK(A) + 1$ ; } else if ( $MARK(channel) == 2$ ) { $MARK(C) = 1$ ; $MARK(A) = MARK(A) + 1$ ; } else if ( $MARK(channel) == 3$ ) { $MARK(A) = 1$ ; $MARK(A) = MARK(A) + 1$ ; }
<i>sense_channel</i>	input	$MARK(channel) == 2 \parallel$ $MARK(channel) == 3$	$MARK(channel) = 0$ ;

Table 1: Gates for Station Submodel

Activity	Distribution Type	Parameter (Rate)
<i>arrival</i>	exponential	varied
<i>access</i>	exponential	10 (normal prio. station) 100 (high prio. station)
<i>finish</i>	exponential	if ( $MARK(channel) == 2$ ) rate = 1 else rate = 5

Table 2: Activity Parameters for Station Submodel



Gate	Type	Enabling Predicate	Function
<i>pda_gate</i>	input	MARK( <i>channel</i> )=1	MARK( <i>channel</i> )=2;

Activity	Distribution Type	Parameter (Rate)
<i>prop_delay_intra</i>	exponential	20

Figure 4: Network Submodel

$(\mathcal{C}_Z, \mathcal{R}_Z)$ ) and a reward structure for the network submodel such that:

$$\mathcal{C}_N(a) = 0 \quad \forall a \in A$$

$$\mathcal{R}_N(\nu) = \begin{cases} 1 & \text{if } \nu = \{(channel, i)\} \\ 0 & \text{otherwise} \end{cases}$$

where  $i$  is the marking of the channel when the action is taking place and  $E[V_{t \rightarrow \infty}]$  is the variable. For example, to determine the fraction of time a propagated packet is on the channel,  $i$  would be taken to be 2. The fractions of time that the bus is idle, an unpropagated message is on the bus, and a corrupted message is on the bus are determined in a similar manner.

SAN-based reward models for the entire network are then constructed by replicating the normal and high-priority stations the required number of times and adjoining them with the network submodel and reward structure. The particular reward structures that are used depend on the particular performance variable being solved. For example, to determine the expected queue length at a station for the homogeneous system (scenario 1), the SAN-based reward model given in Figure 5 can be used. In this figure,  $S_{STATION}$  is the station submodel,  $S_{NETWORK}$  is the network submodel, and  $n$  is the number of stations



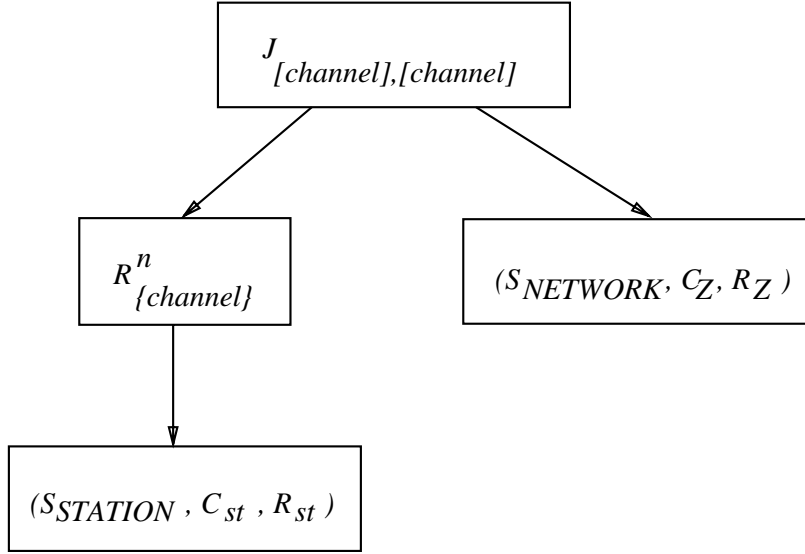


Figure 5: Homogeneous Network

in the network. Similarly, if we define  $S_{HI}$  to be a high-priority station submodel, then a SAN-based reward model which can be used to determine the expected queue length at the high-priority station is given in Figure 6, where  $n$  is now the number of normal priority stations in the model. SAN-based reward models for the other variables can be constructed in a similar manner.

State-level representations for each scenario were constructed using an extension to METASAN<sup>1</sup>[21] that permits the construction and solution of reduced base model representations. Four state space representations were contrasted: the activity-marking state space, the marking state space, and the reduced base model state-space size for each scenario.

The results are presented in Table 3. One can see that for each network configuration considered, the reduced base model construction technique generated significantly fewer states compared to the marking and activity-marking states. Moreover, the rate of growth of the state-space experienced by increasing the number of stations in the network was much smaller for the reduced base model construction technique. In fact, the number of states generated using this technique remained quite small even when the generation of state-spaces using standard techniques became intractable.

The performance of a ten station network was then evaluated using these reduced base

---

<sup>1</sup>METASAN is a registered Trademark of the Industrial Technology Institute.

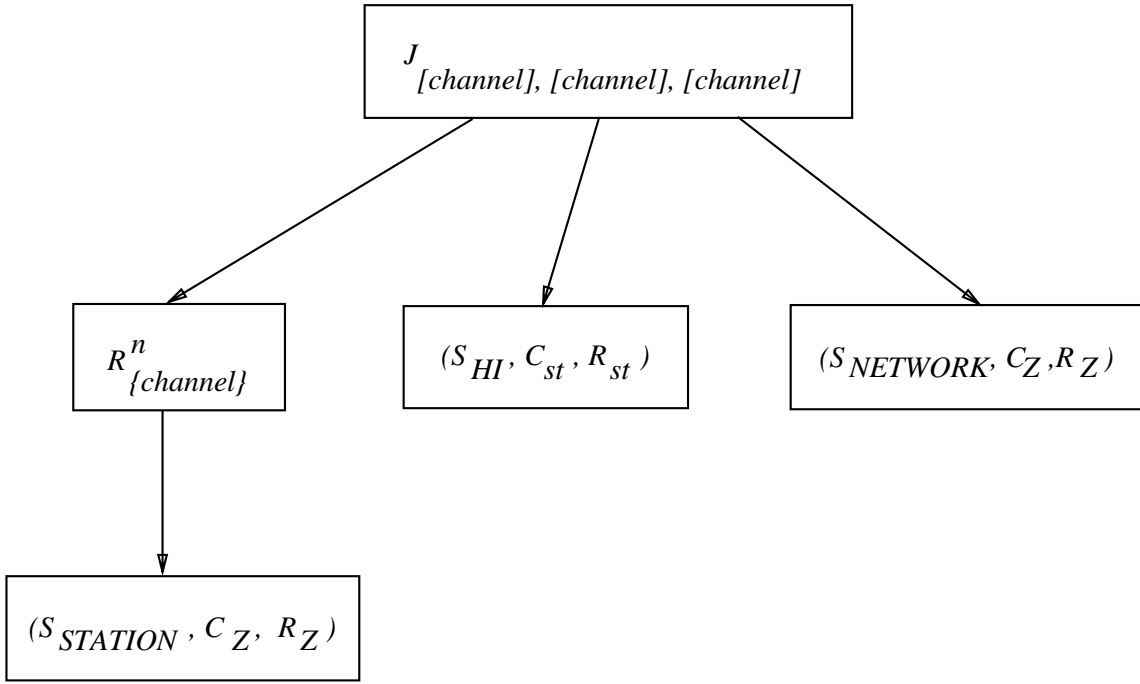


Figure 6: Normal/High Priority Network

Number of Stations	Detailed Base Model		Reduced Base Model	
	<i>am-states</i>	<i>m-states</i>	<i>Scenario 1</i>	<i>Scenario 2</i>
2	135	57	30	57
3	841	261	61	147
4	4277	1041	102	276
5	17820	3873	153	444
6	80000	13883	214	651
7	* <sup>a</sup>	* <sup>a</sup>	285	897
8	* <sup>a</sup>	* <sup>a</sup>	366	1182
9	* <sup>a</sup>	* <sup>a</sup>	457	1506
10	* <sup>a</sup>	* <sup>a</sup>	558	1869
11	* <sup>a</sup>	* <sup>a</sup>	669	2271
12	* <sup>a</sup>	* <sup>a</sup>	790	2712

<sup>a</sup>state-space too large to be computed

Table 3: State-Space Sizes Obtained Using Various Construction Techniques

Fraction of Full Load	Expected Queue Length	Prob. Blocking	Fraction of Time			
			<i>Idle</i>	<i>Unprop.</i>	<i>Prop.</i>	<i>Collision</i>
.10	.0022	.0000	.8957	.0050	.0992	.0002
.20	.0072	.0001	.7939	.0098	.1955	.0009
.30	.0153	.0005	.6963	.0143	.2867	.0026
.40	.0271	.0013	.6046	.0186	.3711	.0057
.50	.0427	.0027	.5202	.0224	.4472	.0103
.60	.0622	.0049	.4438	.0257	.5138	.0166
.70	.0857	.0081	.3762	.0285	.5705	.0248
.80	.1128	.0125	.3174	.0309	.6172	.0346
.90	.1432	.0182	.2670	.0327	.6543	.0459

Table 4: Performance of 10 Station Homogeneous Network

models. Specifically, expected queue lengths and the fractions of time each type of packet was on the bus in steady-state for the network were computed for various network loads. The results of these experiments are presented in Table 4. In this table, the second column is the expected queue length at each station and the third column is the probability that an incoming message is blocked due to a full queue. The remaining columns give the fraction of time the bus is idle, transmitting an unpropagated message, transmitting a propagated message, and clearing a collision. As can be seen by Table 4, the expected queue length grows with increasing load, but remains fairly short even under heavy load conditions.

The second set of experiments investigates the effect of including a high-priority station in the network. The high-priority station is identical to the normal priority station described earlier, except that the mean time between access tries (i.e., one over the rate of activity *access*) is one-tenth that of a normal priority stations. As before, a ten station network is considered, but now nine stations are of the normal priority category and one is a high-priority type. The results of these experiments are given in Table 5. As can be seen by this table, the expected queue length of the high-priority station remains very short even when the network is loaded very heavily. Furthermore, while the other stations in the network experienced some increase in their average queue length (compared to the homogeneous network considered earlier), this increase was slight compared to the decrease in expected queue length at the high-priority station.

## REFERENCES

- [1] P. Heidelberger and S.S. Lavenberg, "Computer performance evaluation methodology",

Fraction of Full Load	Expected Queue Length Normal Prio. Station	Expected Queue Length High Prio. Station
.10	.0022	.0011
.20	.0072	.0043
.30	.0155	.0095
.40	.0275	.0165
.50	.0435	.0251
.60	.0637	.0350
.70	.0878	.0459
.80	.1160	.0575
.90	.1477	.0694

Table 5: Performance of a 10 Station Network with 2 Station Classes

- IEEE Trans. on Computers*, vol. C-33, no. 12, pp. 1195–1220, December 1984.
- [2] L. Kleinrock, *Queueing Systems, Volume I: Theory*, New York: John Wiley, 1975.
- [3] S. Natkin, “Reseaux de Petri stochastiques”, Thèse de Docteur-Ingénieur, CNAM-PARIS, June 1980.
- [4] M. Molloy, *On the integration of delay and throughput measures in distributed processing models*, PhD thesis, UCLA, 1981.
- [5] *Proc. of the First International Workshop on Timed Petri Nets*, IEEE Press, Torino, Italy, July 1985.
- [6] *Proc. of the International Workshop on Petri Nets and Performance Models*, IEEE Press, Madison, WI, Aug. 1987.
- [7] J. F. Meyer, “On evaluating the performability of degradable computing systems”, *IEEE Trans. on Computers*, vol. C-22, pp. 720–731, August 1980.
- [8] M. A. Marsan, G. Balbo, G. Chiola, S. Donatelli, “On the product-form solution of a class of multiple-bus multiprocessor system models”, *The Journal of Systems and Software*, vol. 1, no. 2, pp. 117–124, 1986.
- [9] A. A. Lazar and T. G. Robertazzi, “Markovian Petri net protocols with product form solution”, in *Proc. INFOCOM '87*, San Francisco, CA, Mar. 1987.
- [10] A. Zenie, “Colored stochastic Petri nets”, in *Proc. International Workshop on Timed Petri Nets*, pp. 262–271, Torino, Italy, July 1985.
- [11] C. Lin and D. C. Marinescu, “Stochastic high-level Petri nets and applications”, *IEEE Trans. on Computers*, vol. C-37, no. 7, pp. 815–825, July 1988.
- [12] W. H. Sanders and J. F. Meyer, “Variable driven construction methods for stochastic activity networks,” in *Applied Mathematics and Performance/Reliability Models of Computer/Communication Systems* (ed., G. Iazeolla), North-Holland, 1987.

- [13] A. Movaghar and J. F. Meyer, "Performability modeling with stochastic activity networks," in *Proc. 1984 Real-Time Systems Symp.*, Austin, TX, Dec. 1984.
- [14] J. F. Meyer, A. Movaghar, and W. H. Sanders, "Stochastic activity networks: Structure, behavior, and application," in *Proc. International Workshop on Timed Petri Nets*, Torino, Italy, July 1985, pp. 106–115.
- [15] W. H. Sanders, "Construction and solution of performability models based on stochastic activity networks," Computing Research Laboratory Technical Report CRL-TR-9-88, The University of Michigan, Ann Arbor, MI, August 1988.
- [16] R. A. Howard, *Dynamic Probabilistic Systems, Vol II: Semi-Markov and Decision Processes*, New York: Wiley, 1971.
- [17] W. H. Sanders and J. F. Meyer, "A unified approach for specifying measures of performance, dependability, and performability," to be presented at the *International Working Conference on Dependable Computing for Critical Applications*, Santa Barbara, CA, August 21-23, 1989.
- [18] J. C. Kemeny and J. L. Snell, *Finite Markov Chains*, Princeton: D. Van Nostrand Co., Inc., 1969.
- [19] M. Rosenblatt, *Markov processes. Structure and Asymptotic Behavior*, Berlin: Springer-Verlag, 1971.
- [20] J. L. Peterson, *Petri Net Theory and the Modeling of Systems*, Englewood Cliffs: Prentice-Hall, 1981.
- [21] W. H. Sanders and J. F. Meyer, "METASAN: A performability evaluation tool based on stochastic activity networks", *Proc. ACM-IEEE Comp. Soc. 1986 Fall Joint Comp. Conf.*, Dallas, TX, November 1986.