

# Analysis of Discrete-time Stochastic Petri Nets

W.M.P. van der Aalst<sup>1,2</sup>, K.M. van Hee<sup>1,3</sup>, and H.A. Reijers<sup>1,3</sup>

<sup>1</sup>Eindhoven University of Technology, Department of Mathematics and Computing Science, P.O. Box 513, NL-5600 MB, Eindhoven, The Netherlands

<sup>2</sup>Eindhoven University of Technology, Department of Technology Management, P.O. Box 513, NL-5600 MB, Eindhoven, The Netherlands

<sup>3</sup>Deloitte & Touche Bakkenist, Management and ICT Consultants, P.O. Box 23103, NL-1100 DP, Amsterdam, The Netherlands

corresponding author: H.A. Reijers (hreijers@win.tue.nl)

The Petri net formalism is widely applied in both theoretical and practical settings. For the sake of performance analysis, the original Petri net model has been extended with the notion of time. This paper addresses the different issues involved with this extension. Also, it provides a state-of-the-art overview of different analysis techniques for timed Petri nets. A new analysis technique is presented, which combines the freedom of choosing arbitrary time distributions within a Petri net model on the one hand and efficient computation means on the other.

Key words & Phrases: Petri nets, performance analysis.

## 1 Introduction

We are happy to contribute to the special issue in honor of Jaap Wessels. Directly or indirectly each of us was a scholar of Jaap. Kees van Hee was one of the first Ph.D. students of Jaap (1978). Wil van der Aalst was a Ph.D. student of both Jaap and Kees (1992). Hajo Reijers is currently a Ph.D. student of both Kees and Wil. The three of us are working in Computer Science on the theory and application of Petri nets, while Jaap's research is in stochastic processes. There is a strong relationship between both fields: each Petri net determines the incidence matrix of a discrete, possibly infinite, Markov chain. If random variables are associated to non-deterministic choices in a Petri net, we obtain a Markov process. Computer scientists use Petri nets to model complex systems and to verify if the modeled systems satisfy some correctness criteria (invariant properties and reachability of states). Petri nets are a powerful modeling technique because they provide a way to decompose the states of a system. (Note that a finite Petri net may correspond to an infinite Markov chain.) There is a growing interest in Computer Science to study performance of systems. Stochastic Petri nets extend the traditional Petri net with timing and probability features. As a result, a stochastic Petri net describes a stochastic process. Clearly we did not deviate that much from Jaap's research interests.

The classical Petri net is a directed bipartite graph. The two types of nodes are called *places* and *transitions*. In a Petri net, places and transitions are connected via *arcs*. Places are graphically represented by circles, transitions by boxes or bars. Places can store *tokens*, represented by black dots. A distribution of tokens on the places of a net is called a *marking*, and corresponds to the "state" of the Petri net. A transition of a net is *enabled* at a marking if all its input places (the places from which some edge leads to it) contain at least one token.

An enabled transition can fire: it removes one token from each of the input places, and adds one token to each of its output places. This is called the *firing rule*.

In Figure 1, a Petri net example is depicted. As a matter of coincidence, the Petri net has just as many transitions as places. The places have labels  $p_1, p_2 \dots p_{11}$ ; the transitions  $t_1, t_2 \dots t_{11}$ . In the particular notation that is used for this example, two types of transitions are used. Transitions that are drawn as bars are timeless when fired; the transitions that are depicted as boxes consume time. As will be discussed in the next section, this is one of the many ways to add time behavior to Petri nets. Note that both types of transitions respect the firing rule. In the Petri net depicted in Figure 1, places  $p_8$  and  $p_{10}$  are marked. The transitions that are enabled are  $t_7$  and  $t_{10}$ . In particular, transition  $t_{11}$  is *not* enabled. If transition  $t_{10}$  would fire on basis of the token in  $p_8$ , then  $t_{11}$  becomes enabled as a token would be added to  $p_9$ . As a final remark, transitions  $t_7$  and  $t_{10}$  are said to be *in conflict*, as they compete for the same token to fire.




Fig. 1 A Petri net.

This paper focuses on *timed Petri nets*. Because there are many ways to incorporate time in Petri nets, we present in Section 2 an up-to-date overview of the different timing mechanisms. In Section 3, we discuss the state-of-the-art analysis techniques for the various timed Petri net models. Based on this discussion, we motivate the analysis technique presented in this paper in Section 4. The technique differs from existing techniques in the sense that it combines the use of arbitrary stochastic timing with efficient means of computation and the rich expressiveness of the Petri net model.

## 2 Time in Petri nets

Petri nets were originally proposed as a causal model without any notion of time or probability. In fact, the concept of time was intentionally avoided in the original work by Carl Adam Petri (1962). The addition of time restricts the behavior of the net, i.e., the dynamic behavior of the Petri net is only partially reflected by the network structure. However, for many practical applications, the addition of time is a necessity. Without an explicit notion of time it is not possible to analyze the performance (e.g., throughput and utilization) of the modeled system. Since the early seventies there has been a discussion within the Petri net community on the addition of time. More theory-oriented researchers oppose or simply ignore timing issues. More application-oriented researchers advocate and investigate different timing mechanisms and analysis techniques. We typically use a mixed approach, i.e., we start with a Petri net extended with time. In fact, we often use high-level Petri nets extended with color (data), time, and hierarchy as described by Van der Aalst (1993), Chiola et al. (1990), Ghezzi et al. (1991), Van Hee (1994), Jensen (1996), Lin and Marinescu (1987), Morasca et al. (1991), and Zenie (1985). For analysis purposes we abstract from some of these extensions. The reason for such abstractions is the need to make analysis tractable: Many properties are undecidable for Petri nets extended with data and/or time. In the next paragraphs, we will discuss the different aspects of incorporating time in Petri nets.

## 2.1 Location of delay

There are many ways to introduce time into Petri net models. As explained in the introduction, a Petri net consists of places and transitions connected via arcs. Therefore, time can be associated with *places*, *transitions*, or *arcs*. In most timed Petri net models, transitions determine time delays. In only a few models, time delays are determined by places and/or arcs. It seems that it is more natural to associate time to transitions: Transitions represent activities and activities take time. However, authors such as Sifakis (1977, 1980) argue that it is more convenient to associate time to places since this leaves the original firing rule intact, i.e., enabling and firing are instantaneously. For high-level Petri nets with colored tokens (i.e., tokens carry a data value), it is most natural to attach timestamps to tokens, as done by Van der Aalst (1993), Van Hee (1994), and Jensen (1996). The timestamp indicates the time a token becomes available for consumption. In the models presented by Van der Aalst (1993), Van Hee (1994), and Jensen (1996) transitions set the timestamps of produced tokens, i.e., time delays are determined by transitions.

## 2.2 Type of delay

Independent of the choice where to put the delay (i.e., transitions, places, or arcs), several types of delays can be distinguished. In this paper, we distinguish between *deterministic*, *non-deterministic*, and *stochastic* delays. Many of the older timed Petri net models by authors such as Van Hee et al. (1989), Ramchandani (1984), Sifakis (1977), Wong et al. (1985), and Zuberek (1980) use deterministic delays, i.e., the delay assigned by a transition, place, or arc is fixed. Deterministic delays allow for simple analysis methods but have limited applicability. In real applications, delays correspond to the duration of activities which are typically variable. Therefore, fixed delays are often less appropriate. There are two ways to describe the intrinsic variability. One way is to assume constraints on delays (e.g., it takes less than 15 minutes to type a letter), i.e., non-deterministic delays. Another way is to assume a probability distribution for each delay, i.e., stochastic delays. Most of the models handling non-deterministic delays use time intervals to specify the duration of the delay. Merlin (1974, 1976) introduced such a model in the early seventies. Other models using interval timing have been proposed by Van der Aalst (1993, 1994), Van der Aalst and Odijk (1995), Berthomieu and Diaz (1991), and Berthomieu and Menasche (1993). However, most of the timed Petri net models use stochastic delays. In these models each delay is described by a probability distribution. To make analysis tractable typically only a restricted set of probability distributions is allowed. In the SPN model by Florin and Natkin (1982), only exponential delays (i.e., delays described by a negative exponential probability density function) are allowed. The widely used GSPN model by Marsan et al. (1984) allows for both immediate transitions (i.e., transitions with no delay) and timed transitions (i.e., transitions with exponential delays). Models allowing for arbitrary probability distributions typically defy exact analytical analysis. Another way to classify the types of delay used in a timed Petri net model is to distinguish between *discrete* and *continuous* delays. Most discrete models use the natural numbers as the time domain. Continuous models typically use the set of non-negative real numbers as the time domain. Nearly all timed Petri nets allow for continuous delays. The analysis method described in the second part of the paper assumes discrete delays. This is not a real limitation since a fine-grained discrete time domain can be used to approximate a continuous time domain.

## 2.3 Preselection versus race semantics

Adding time to Petri nets requires a redefinition of the enabling and firing rules. In a classical Petri net a transition is enabled if each of the input places contains enough tokens (typically

1), only enabled transitions can fire, and firing is instantaneously (i.e., the moment tokens are consumed from the input places, tokens are added to the output places). Transitions are in conflict if they share input places. Note that firing a transition in conflict with other transitions may disable some or all of these transitions. The choice between several enabled transitions in conflict with each other is resolved in a non-deterministic manner. When adding time to a Petri net the enabling and firing rules need to be modified to specify how conflicts are resolved (i.e., the relation between enabling and firing) and whether firing is instantaneous or not (the semantics of the firing rule). Clearly, these two issues are related. Assume that transitions determine the delays. If firing is instantaneous (i.e., does not take any time), then it is necessary to associate time to the enabling of a transition. If time is associated to the enabling, there is no need to explicitly define how conflicts are resolved, i.e., enabled transitions “race” against each other and the one that is scheduled to fire first will fire. This firing/enabling semantics is called the *race semantics*. It is also possible to specify the way conflicts are resolved explicitly. This firing/enabling semantics is called the *preselection semantics*. For example, priorities or probabilities are used to resolve conflicts. In the preselection semantics there is no race between enabled transitions: The moment transitions become enabled one of the enabled transitions is selected. Race semantics are typically combined with instantaneous firing, i.e., time is in the enabling of transitions. Therefore, we also use the term *enabling delays* to refer to these semantics. Preselection semantics are typically combined with *holding times*, i.e., tokens reside for some time inside a place or transition. Note that for race semantics the resolution of conflicts and the delay are handled by the same mechanism. For preselection semantics the mechanism to resolve conflicts is separated from the actual delay. Most of the stochastic Petri nets by authors such as Balbo and Silva (1998), Florin and Natkin (1982), Ajmone Marsan (1990), and M. Ajmone Marsan et al. (1984, 1985, 1986, 1995) use race semantics. As established by Van der Aalst (1992), race semantics allow for a more direct translation to Markov chains and timed Petri nets using race semantics are more expressive than timed Petri nets using preselection semantics. For example, race semantics allow for a compact representation of time-outs. Preselection semantics are more intuitive and easier to use. Therefore, most of the high-level Petri nets such as the ones described by Van der Aalst (1993), Van Hee (1994) and Jensen (1996) support preselection semantics. Other authors such as Razouk and Phelps (1984) propose a mixture of race and preselection semantics.

For preselection semantics the delays (i.e., holding times) can be associated to the firing of a transition (e.g. described by Berthomieu and Diaz (1991)) or the minimal time a token spends in a place (e.g. described by Sifakis (1980)).

For race semantics the delays are associated to the enabling time. Note that an enabled transition can be disabled by another transition in case of a conflict. Such a transition loses the race and will not fire. If the transition becomes enabled again, a new race starts. In this new race there are several possibilities for the new enabling time of this transition. Authors such as Balbo and Silva (1984), Ajmone Marsan et al. (1985, 1995) typically distinguish three so-called memory policies: *age memory*, *enabling memory*, and *reset memory*. For age memory, the remaining enabling time is frozen the moment the transition becomes disabled and is resumed the moment the transition becomes enabled again. For enabling memory, a new enabling time is sampled every time a transition becomes enabled, i.e., previously interrupted transitions have to start from scratch. For reset memory, a new enabling time is sampled every time a transition fires, i.e., also transitions not in conflict with the transition that fired are interrupted and have to resample a new enabling time. It is interesting to note that for stochastic Petri nets with just exponential delays the three memory policies coincide.

The memoryless property of the negative exponential probability density function makes the residual enabling time statistically equivalent to the originally sampled enabling time.

#### 2.4 Capacity, priority, and queuing policy

For timed Petri nets, the capacity of places and transitions is relevant. Places can have a limited capacity to restrict the number of tokens residing in a place at the same moment in time. Transitions can have a capacity to limit the number of concurrent enablings/firings of the same transition. Consider a transition with one input place containing three tokens. Is this transition enabled three times under race semantics? Can the transition fire concurrent with itself under preselection semantics? To answer these questions, we identify three types of capacity related semantics: *single server semantics*, *multiple server semantics*, *infinite server semantics*. For single server semantics the capacity of a place/transition is 1, for multiple server semantics the capacity of a place/transition is some integer  $k$ , and for infinite server semantics there are no capacity restrictions. Most timed Petri net models assume infinite server semantics.

Several timed net models allow for a priority mechanism, i.e., if multiple transitions compete for the same token, the transition with the highest priority fires. Note that the priority mechanism can be used for preselection purposes. In the widely used GSPN model by Marsan et al. (1984) immediate transitions (i.e., transitions with no delay) have priority over timed transitions.

Some Petri net models allow for the specification of queuing policies. However, since tokens in the same place (of the same color) are indistinguishable, it often does not make any sense to choose a queuing discipline. In general priorities (i.e., not transition priorities but token priorities), random selection (RS), and processor sharing (PS) are easy to handle in a stochastic Petri net, as established by Balbo and Silva (1984). State-dependent queuing disciplines such as first-come-first-served, last-come-first-served, longest-job-first, and earliest-due-date-first are more difficult to represent and analyze.

In this section we presented an overview of the various ways time can be incorporated in Petri nets. In the second part of this paper, we will use a timed Petri net model with time in transitions using preselection semantics, i.e., firing a transition takes time and conflicts are resolved independent of the firing time using probabilities. The type of delay is stochastic. However, unlike most stochastic Petri nets, we use a discrete time domain. On this discrete time domain, we allow arbitrary, but finite, probability distributions. The model used in this paper does not support priorities. However, transition priorities could easily be added. Since the Petri nets considered in this paper are safe (i.e., the maximum number of tokens in a place is one), capacity and queuing policy are not relevant.

### 3 Analysis of timed nets

As indicated in the previous section, there are many ways to introduce time in Petri nets. All timed Petri net models are executable, i.e., it is possible to construct a trace of the modeled system by playing the token game. Therefore, simulation can be used to analyze the model. If all non-determinism is replaced by stochastic measures (i.e., delays and conflict resolution), then simulation can be used to obtain confidence intervals for performance measures such as utilization and throughput. Because simulation does not require difficult mathematical techniques, it is easy to understand for people with a non-technical background. Simulation is also a very flexible analysis technique, since it does not set additional constraints. However,

sometimes simulation is expensive in terms of the computer time necessary to obtain reliable results. Another drawback is the fact that (in general) it is not possible to use simulation to *prove* that the modeled system has the desired set of properties. In the remainder, we discuss alternative analysis techniques to overcome the limitations of simulation approaches. Since analysis techniques are typically restricted by the type of delay, we first consider timed Petri nets with deterministic timing, then timed Petri nets with non-deterministic timing, and finally timed Petri nets with stochastic timing.

### 3.1 *Deterministic timing*

There are several methods to calculate upper and lower bounds for the *cycle time* of a timed Petri net with deterministic delays. For example, check authors such as Murata (1992), Ramamoorthy and Ho (1980), Ramchandani (1984), and Sifakis (1939). The cycle time is a criterion for the performance of the system. For a specific class of deterministic timed Petri nets, the so-called Timed Event Graphs, the exact cycle time can be computed quite efficiently, see Ramamoorthy and Ho (1980) and Chretienne (1983). Other researchers such as Zuberek (1980) analyze deterministic timed Petri nets by building the reachability graph. Although this requires a lot of computing effort, such a graph can be used to answer a variety of questions. A serious drawback of these methods is the fact that in many real systems the activity durations are not fixed, i.e., they vary because of disturbances and other interferences. Assuming deterministic delays often results in inaccurate results.

### 3.2 *Non-deterministic timing*

Most timed Petri nets models using non-deterministic delays, such as described by Van der Aalst (1993, 1994), Van der Aalst and Odijk (1995), Berthomieu and Diaz (1991), Berthomieu and Menasche (1993), Merlin (1974), and Merlin and Faber (1976), use intervals to describe lower bounds and upper bounds for the duration of activities. The method presented by Berthomieu, Diaz and Menasche (1983, 1991) uses Merlin's (1974) timed Petri net model. The method generates a reachability graph where nodes represent state classes instead of states. Sets of linear equations are solved to calculate these state classes. The method is able to reduce the number of states by using a relative time scale. Another method using interval timed colored Petri nets was presented by Van der Aalst (1993). This method uses an absolute time scale and allows for colored tokens. The method also generates a reachability graph where nodes represent state classes. The number of states is reduced by exploiting "timed" specialization and generalization properties. Van der Aalst and Odijk (1995) describe an application of this method and Van der Aalst (1992) gives two additional analysis methods based on interval timing.

### 3.3 *Stochastic timing*

The majority of stochastic Petri nets models uses a continuous time domain. In these models, each delay is described by a probability density function. For arbitrary probability density functions, only simulation or approximation are feasible analysis techniques. Therefore, many stochastic Petri nets models impose restrictions on the type of delay distribution. In the SPN (Stochastic Petri Net) model as described by Florin and S. Natkin (1982) and Molloy (1981) only exponential delays are allowed. Due to the memoryless property of the exponential distribution and the race semantics, Florin and S. Natkin (1982) and Molloy (1981) show that SPN's are isomorphic to continuous time Markov chains. The number of states of the Markov chain corresponds to the number of reachable markings of the SPN. The GSPN (Generalized Stochastic Petri Net) model extends the SPN model with immediate transitions. Immediate transitions fire without any enabling time and have priority over timed

transitions (i.e., transitions with exponential enabling times under the race semantics). A marking is vanishing if an immediate transition is enabled. A marking is tangible if only timed transitions are enabled. The GSPN model distinguishes between these two types of markings: only tangible markings consume time, i.e., the average sojourn time of vanishing states is zero and the average sojourn time of tangible states is positive. The dynamics of a GSPN corresponds to a semi-Markov process: The embedded Markov chain which ignores the sojourn time in each state is a discrete time Markov chain. By using the embedded Markov chain, it is fairly straightforward to calculate various performance measures. Note that only the tangible states consume time. Therefore, the vanishing markings are not relevant for most performance measures. Therefore, as shown by Balbo and Silva (1984) and Ajmone Marsan et al. (1985, 1995), it is possible to reduce the number of states by eliminating the vanishing markings in the embedded Markov chain.

The GSPN model has been extended in various directions. First of all, the GSPN model has been extended with marking dependent transition probabilities and enabling delays. It is easy to see that such an extension can be handled by using an embedded Markov chain as long as immediate and timed transitions do not interfere. Second, the GSPN model has been extended to allow for other types of delay distributions (i.e., non-exponential). Basically there are two ways to incorporate non-exponential delays. First of all, it is possible to introduce transitions with arbitrary delay distributions as long as none of these transitions can be enabled concurrently. The work of Ajmone Marsan and Chiola (1987) is an example of the DSPN model which allows for timed transitions which have either fixed (i.e., deterministic) or exponential enabling times. The DSPN model can be analyzed as a semi-Markov process as long as only one deterministic transition is enabled at the same time and the enabling memory policy is assumed. Several variations and refinements of the DSPN have been proposed in literature (Balbo and M. Silva (1998) give pointers). Another approach to incorporate non-exponential delays is to allow for delay distributions which can be represented by a continuous time Markov chain. Examples of such delays are the Erlang, the hyperexponential, and the phase-type distribution. The possibility to incorporate such delays was already mentioned by Florin and Natkin (1982) and Molloy (1981). The relation between the various memory policies and phase-type distributed transitions is discussed by Balbo and Silva (1984) and Ajmone Marsan et al. (1995). Using non-exponential delays which are expanded to multiple phases in the corresponding Markov chain typically results in Markov chains which are difficult to analyze. (In worst case, the size of the Markov chain grows exponentially in the number of phases.)

Nearly all stochastic Petri net models described in literature use a continuous time domain and are analyzed using Markov-chain analysis. For most applications, the GSPN model and its extensions are a convenient and practical tool. However, these models have two potential drawbacks. First of all, the size of the Markov chain typically grows exponentially in the size of the corresponding Petri net. Second, the use of non-exponential delays is restricted: Only by restricting the topology of the Petri net or by expanding the Markov chain to encode the phases of a phase-type distribution it is possible to allow for non-exponential delays. Net-driven decomposition techniques can be used to partially alleviate the state-explosion problem, as shown by Balbo and Silva (1984). Unfortunately, it is almost impossible to handle arbitrary delay distributions. Therefore, we propose an alternative analysis technique. The technique uses a discrete time domain instead of a continuous time domain. An arbitrary probability distribution can be used for specifying the firing delays. (Recall that the stochastic Petri net used in this paper uses preselection and holding durations for transition firings.) The constraints on the topology are quite acceptable: Typical routing constructs such as

sequencing, parallelism, choice, and iteration are supported. Since the method allows for arbitrary distributions, the algorithm is potentially inefficient, i.e., distributions over the discrete time domain are represented explicitly. However, the Fast Fourier Transform is used to improve the efficiency of several analysis routines. Moreover, the method allows for efficient construction rules and decomposition techniques.

#### 4 Discrete time stochastic nets

The analysis method presented in this section is applicable to a subclass of the so-called *discrete time nets*. As a matter of fact, this subclass coincides with the process-algebraic expressions that can be constituted with the merge ( $\parallel$ ), choice ( $+$ ), sequential ( $\cdot$ ), and star ( $*$ ) operators in the process algebra ACP, as described by Bergstra and Klop (1984). From a practical viewpoint, it is possible to model the majority of typical business processes as found in organizations as banks, insurance companies, and governmental agencies with this particular subclass.

The nets under consideration in this section are composed out of smaller blocks. In the next paragraph, we will start to describe the composition method (synthesis). Next, the notion of throughput is introduced, so that the computation of the throughput time can be explained. An example is added to illustrate the approach. Finally, possible extensions of the approach are discussed in this section.

##### 4.1 Synthesis

The models on which we will apply the analysis technique are based on the classical Petri net.

**Definition (Petri net)** A *Petri net* is a triplet  $(P, T, F)$ :

- $P$  is a finite set of places,
- $T$  is a finite set of transitions ( $P \cap T = \emptyset$ ),
- $F \subseteq (P \times T) \cup (T \times P)$  is a set of arcs (flow relation).

A place is called an *input place* of a transition  $t$  iff there exists a directed arc from  $p$  to  $t$ . Place  $p$  is called an *output place* of transition  $t$  iff there exists a directed arc from  $t$  to  $p$ . We use  $\bullet t$  to denote the set of input places for transition  $t$ . The notations  $t\bullet$ ,  $\bullet p$  and  $p\bullet$  have similar meanings, e.g.  $p\bullet$  is the set of transitions sharing  $p$  as an input place.

Using the notion of a classical Petri net, we can introduce the discrete time stochastic nets. Our approach is applicable to a sub-class of these nets (cf. Van der Aalst (1998)).

**Definition (Discrete time stochastic nets, dts-net)** A Petri net  $PN = (P, T, F)$  is a *discrete time stochastic net*, or *dts-net* for short, iff:

1.  $PN$  has two special places  $i$  and  $o$ ; place  $i$  is a source place:  $\bullet i = \emptyset$ ; place  $o$  is a sink place:  $o\bullet = \emptyset$ ,
2. if transition  $t^*$  would be added to the set of transitions  $T$  and the arcs  $(o, t^*)$  and  $(t^*, i)$  would be added to the flow relation  $F$  of  $PN$ , for every two places (transitions)  $x$  and  $y$ , there is a directed path leading from  $x$  to  $y$ .

The first requirement in the definition of a dts-net reflects the beginning and termination of the process. That event is to be handled in such a way that a desired end situation is reached. The second requirement in the definition ensures that there are no dangling transitions or



places. In other words, transitions or places are not permitted to be part of a Petri net model if they do not contribute to the dynamic behavior of the model. A dts-net specifies the dynamic behavior of a *single case in isolation*. Note that the net in Figure 1 is a dts-net.

The synthesis method to construct a dts-net is comparable with the top-down synthesis technique as used by Valette (1979) and Van der Aalst (2000) in which a transition is substituted by a block. We will use four blocks that can be used to substitute transitions. These are: sequence, choice, parallelism, and iteration. The blocks that express these constructs are depicted in Figure 2.




Fig. 2. Sequence (a), iteration (b), parallelism (c), and choice (d).

Arc labels occur in the (b) iteration and (d) choice blocks. They represent the values of a Bernoulli-distributed random variable that is associated with these blocks. An independent draw from its distribution function determines the route of the flow of tokens through the net. In other words, conflicts in the blocks are resolved using preselection semantics. Each new application of a block is accompanied by the introduction of a new, independent random variable. Furthermore, two types of transitions are distinguished. *Immediate* or *timeless* transitions are depicted in the blocks as black bars; the rectangular blocks are transitions to which transitions delays (i.e., holding times) are associated for the firing of a transition.

As the starting point of each Petri net model construction we will take a simple start net. This net is depicted in Figure 3.




Fig. 3. The start net SN.

The net consists of one transition, a source place, a sink place, relations between them, and an initial marking of the source place. It is easy to see that this Petri net is a dts-net. We will refer to this specific Petri net as *SN*, for start net. Net synthesis, then, is straightforward. The timed transitions in the start-net can be substituted by one of the blocks, after which in a recursive fashion each of the resulting timed transitions may again be replaced by a block. For a formal description of this substitution, the reader is referred to Valette (1979). Note that the nets that can be constructed in the described fashion coincide with the process algebraic expressions. A net synthesis of the dts-net as introduced in Figure 1, is depicted in Figure 4. As the start net is a dts-net, it is easy to verify that the resulting net is a dts-net too.




Fig. 4. Synthesis of a dts-net from the start net.

The delay or holding time of each transition firing is called the *service time*. As stated before, in our approach a delay characteristic is associated with each timed transition. All service times for one specific timed transition are independently sampled on basis of this probability distribution. We will call this distribution the *service distribution*. Its matching probability density is the *service density*.

**Definition (service time, service density, service distribution, upper bound)** The time which is taken by a service of transition  $t \in T$  within a dts-net  $W = (P, T, F)$  is called the service time. The service time is a discrete random variable  $\underline{t}$ . Its matching probability density  $f_t : \mathbb{N} \rightarrow \mathbb{R}$  is called the service density ( $\mathbb{N}$  is the set of natural numbers and  $\mathbb{R}$  is the set of non-negative reals):

$$f_t(k) = \mathbb{P}(\underline{t} = k), \text{ for } k \in \mathbb{N}.$$

Its matching probability distribution  $F_t : \mathbb{N} \rightarrow \mathbb{R}$ , is called the service distribution:

$$F_t(k) = \mathbb{P}(\underline{t} \leq k), \text{ for } k \in \mathbb{N}.$$

The service time  $\underline{t}$  is bounded: there is an upper bound  $u_t \in \mathbb{N}$  which is the smallest value such that for all  $j \in \mathbb{N}$  and  $j \geq u_t$  holds that  $f_t(j) = 0$ .

Our assumption of the service time to be discrete is no real constraint: for practical purposes it is always possible to find an appropriate representation. As we will see, we do need the boundedness of the service time to perform some of the computations to come.

For each dts-net the *throughput time* is now defined as the time that elapses between the arrival of a token at the source place and the corresponding arrival of a token in the sink place. Similar to the service time notions, it is possible to distinguish the throughput distribution and throughput density of a dts-net. Assuming the service densities to be known of each of the transitions within a block, we will show how the throughput density of an entire block can be computed. Each of the blocks requires a specific algorithmic approach. For the computations to come, we assume a source and a sink place to be added to each of the blocks, so that our notion of throughput time is applicable to these blocks.

**Sequence.** Consider the sequence block B in Figure 2 with two transitions  $t$  and  $u$ . We want to compute throughput density  $f_B$ , given  $f_t$  and  $f_u$ .

$$\text{Let } y \in \mathbb{N}, f_B(y) = \sum_{i=0}^y \mathbb{P}(\underline{t} = i \wedge \underline{u} = y - i) = \sum_{i=0}^y \mathbb{P}(\underline{t} = i) \mathbb{P}(\underline{u} = y - i) = f_t \otimes f_u(y).$$

To constrain the computation effort drastically, the convolution  $f_t \otimes f_u$  can be computed with the Fast Fourier Transform and its inverse. Therefore, we can compute a vector representation of  $f_t \otimes f_u$  in  $\theta(n \log n)$  time, with  $n$  the smallest power of two that is at least twice as large as the maximum of the upper bounds of tasks  $s$  and  $t$ .

**Iteration.** Now consider the iteration block B as depicted in Figure 2. We want to compute throughput density  $f_B$ , given  $f_s$  and  $f_t$ .

$$\text{Let } y \in \mathbb{N}, \text{ then } f_B(y) = \sum_{n=0}^{\infty} (1-a) a^n \mathbb{P} \left( \sum_{j=1}^{n+1} t_j + \sum_{j=1}^n u_j = y \right) = \sum_{n=0}^{\infty} (1-a) a^n \mathbb{P} \left( \bigotimes_{j=1}^{n+1} f_t \otimes \bigotimes_{j=1}^n f_u(y) \right),$$

$$\text{with notation } \bigotimes_{j=1}^n a_j = a_1 \otimes a_2 \dots \otimes a_n.$$

Now consider  $\bar{B}$  - the vector representation of  $f_B$ . Its Discrete Fourier Transform is:

$$DFT_l(\bar{B}) = DFT_l \left( \sum_{n=0}^{\infty} (1-a) a^n \left( \bigotimes_{j=1}^{n+1} \bar{t} \otimes \bigotimes_{j=1}^n \bar{u} \right) \right) = \sum_{n=0}^{\infty} (1-a) a^n DFT_l^{n+1}(\bar{t}) DFT_l^n(\bar{u}) = \frac{(1-a) DFT_l(\bar{t})}{1 - a DFT_l(\bar{t}) DFT_l(\bar{u})},$$

with yet unknown  $l$ .

We do not know yet an approximated, proper size  $l$  of the vectors we have to “feed” the *DFT*. We cannot expect  $f_B$  to have an exact upper bound. After all,  $t$  and  $u$  can be executed infinitely often if  $\alpha$  is non-zero. We will show how an approximated, relevant length of  $\bar{B}$  can be determined before actually computing  $\bar{B}$ . We would be happy to find a value  $v$  such that for some very small  $\epsilon$  holds that  $\mathbb{P}(\underline{B} \geq v) \leq \epsilon$ . Using Chebyshev’s inequality - for any random variable  $\underline{x}$  for which  $E\underline{x}^2$  exists holds that  $\mathbb{P}(|\underline{x} - E\underline{x}| \geq c) \leq \frac{\text{var } \underline{x}}{c^2}$  - and the mean and variance of the throughput density  $f_B$ , it can be determined which probability part of the density falls before or after a hypothetical border. As the service time for any transition  $t$  is denoted by  $\underline{t}$  we will denote its mean by  $E\underline{t}$  and its variance by  $\text{var } \underline{t}$ . The mean and variance of  $\underline{B}$  are:

$$E\underline{B} = \frac{E\underline{t} + aE\underline{u}}{1-a}, \quad \text{var } \underline{B} = \frac{\text{var } \underline{t} + a \text{var } \underline{u}}{1-a} + a \left( \frac{E\underline{u} + E\underline{t}}{1-a} \right)^2,$$

so that we can derive using Chebyshev and our desired value  $\epsilon$  that:

$$v \geq E\underline{B} + \sqrt{\frac{\text{var } \underline{B}}{\epsilon}}.$$

Concluding, given  $f_t$  and  $f_u$ , we can compute a vector representation of  $f_B$  for the iteration block by using the *DFT*:

$$DFT_v(\bar{B}) = \frac{(1-a) DFT_v(\bar{t})}{1 - a DFT_v(\bar{t}) DFT_v(\bar{u})}, \quad \text{with } v \text{ the smallest power of two such that } v \geq E\underline{B} + \sqrt{\frac{\text{var } \underline{B}}{\epsilon}}.$$

With the *DFT* we can compute a vector representation of  $f_B$  in  $(v \log v)$  time, with  $v$  as specified. To appreciate its efficiency we have to establish the computing time of calculating  $f_B$  in a straightforward manner. The complexity of this calculation depends on the maximal

number of successive times that transitions  $t$  and  $u$  can be executed. We know that if both  $f_t(0)$  and  $f_u(0)$  are equal to zero, at most  $\nu$  executions of these transitions are of interest. Any more executions of transitions  $t$  and  $u$  would result in throughput times that we do not take into consideration. As a result, a straightforward approach requires the convolution of  $\nu$  times the function  $f_t$  and  $f_u$ . This is an operation requiring  $\theta(n^\nu)$  time, with  $n$  the maximum of upper bounds of transitions  $t$  and  $u$ . A comparison with the  $\theta(\nu \log \nu)$  time required by our newly found computation method illustrates the efficiency of the latter.

**Parallelism.** The block we will consider next is the parallel block. Consider the parallel block  $B$  in Figure 2 with transitions  $t$  and  $u$ . Due to the structure of  $B$ , transitions  $t$  and  $u$  can be executed in parallel. We want to compute throughput density  $f_B$ , given  $f_t$  and  $f_u$ .

$$\begin{aligned} \text{Let } y \in \mathbb{N}, f_B(y) &= \mathbb{P}(\underline{t} \max \underline{u} = y) = f_t(y) \sum_{i=0}^y f_u(i) + f_u(y) \sum_{j=0}^{y-1} f_t(j) \\ &= \begin{cases} f_t(y)F_u(y) + f_u(y)F_t(y-1), & y > 0 \\ f_t(y)f_u(y), & y = 0 \end{cases} \end{aligned}$$

The computation of the distribution function  $F_t$  can be performed in  $u_t$  steps, just as the distribution function  $F_u$  can be done in  $u_u$  steps. Therefore, the total computation of  $f_B$  can be done in  $\theta(t)$  time, with  $t$  equal to the maximum of  $u_t$  and  $u_u$ .

**Choice.** The final block we will consider in this section is the choice block. Initiating choice block  $B$  in Figure 2 results in either the execution of transition  $t$  or transition  $u$ , with respective probabilities  $\alpha$  and  $1 - \alpha$ . We would like to compute throughput density  $f_B$ , given  $f_t$  and  $f_u$ .

$$\text{Let } y \in \mathbb{N}, f_B(y) = \alpha f_t(y) + (1 - \alpha) f_u(y)$$

From this expression follows that we can compute  $f_B$  in  $\theta(n)$  time, with  $n$  equal to the maximum of  $u_s$  and  $u_t$ .

## 4.2 Computation

Now that for each block it is clear how the overall service time characteristics can be computed, the computation of the throughput time of a complete synthesized dts-net is rather straightforward. If the reverse direction is taken of the synthesis route, the service time characteristic for each specialized transition can be computed, which - in its turn - can be used to compute the service characteristic of a transition that has been specialized earlier during the synthesis. Finally, the service time characteristic of the transition in the start net can be computed. As a result, the throughput time characteristic of the complete dts-net is known.

To illustrate the presented algorithms we have computed the throughput density of the synthesized dts-net depicted in Figure 5. Each of the five timed transitions in the dts-net has been modeled to behave in accordance with a distinct, quirky service density. These service densities, together with the dts-net under consideration, are depicted in Figure 5. The service densities in this example are bounded by 64 time units.

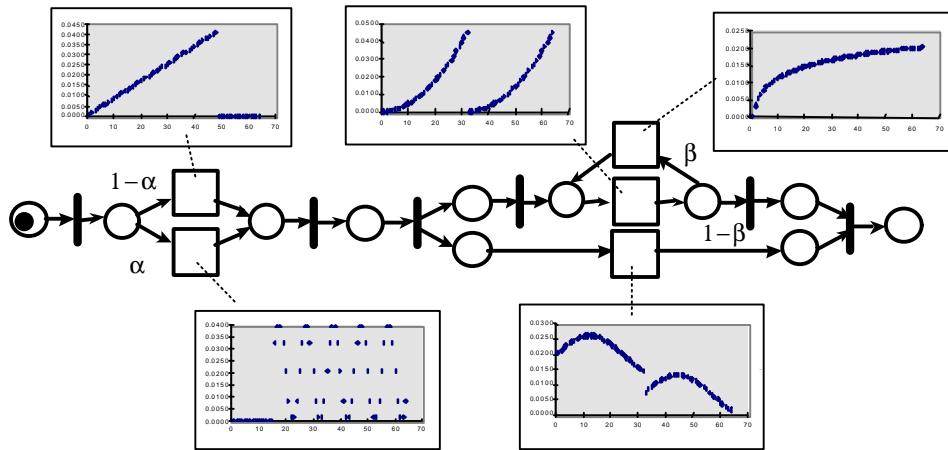


Fig. 5. Dts-net with service time distributions.

The probabilities  $\alpha$  and  $\beta$  of the model have been set on 0.3 and 0.15 respectively. For the computation of the throughput density of the iteration block an inaccuracy ( $\epsilon$ ) of 1 percent has been allowed. The resulting throughput density for the entire dts-net is depicted in Figure 6.




Fig. 6. Resulting throughput time density

### 4.3 Extensions of the approach

The approach that is presented in this section can be extended in at least two directions. The first and most straightforward direction is to add more blocks to the current set of synthesis blocks. The subclass of dts-nets that can be synthesized thus becomes larger, and the application domain grows. At the time of writing, the authors are working on several complex blocks to be added.

We will sketch another interesting direction of extension informally. Intuitively, it seems possible to compute the throughput time of each free-choice, acyclic and sound dts-net. A dts-net is free-choice iff, for every two places  $p$  and  $q$  either  $(p\bullet \cap q\bullet) = \emptyset$  or  $p\bullet = q\bullet$ . A dts-net is acyclic if there is no directed path within the net leading from either a place or a transition to itself. Lastly, a dts-net is sound if it is ensured that a token in the source place will always lead to a token in the sink place, while the rest of the net is empty. On basis of Petri net theory it can be proven that, starting with one token in the source place, each place will contain at most one token during execution. Moreover, the outcome of each conflict can be modeled as a random variable that is independent of the execution order of the net. As a result, it is possible to define an arrival function for each single place in the net. That function expresses the time that elapses between the marking of the source net and the arrival of a token in the particular place. Of course, the arrival function of the sink place is equal to the throughput distribution of the net. Note that this informally sketched method is stochastically feasible on basis of formal reasoning about *the structure and behavior of the net*.

If we mix both approaches - the one presented in this paper and the one described above – we have a powerful toolbox to analyze hybrid nets. Suppose we have a hierarchical dts-net, such that each transition in the net (recursively) can be a larger Petri net in itself. Considered on the highest level, a dts-net may be free-choice, acyclic and sound, but it cannot be constructed by our block synthesis. On a second level, the Petri nets that determine the behavior of the transitions may be synthesized out of blocks including cyclic constructions such as the iteration block. Yet a following, deeper level in the nets may be free-choice, sound and acyclic again or constructed from blocks, etc. On each level of the net we can apply the feasible approach, until the throughput behavior of the entire net is determined. This hybrid approach is illustrated in Figure 7.




Fig. 7. A hybrid approach.

The method works, as the transitions on each level of consideration are “black boxes” for the analysis approach that is applied. As a result of the mixed approach, the analyzable class of dts-nets has grown significantly. In a forthcoming paper, we will describe this approach in detail.

### Acknowledgement

The authors are grateful to Fred Steutel, Eindhoven University of Technology, for his hint to use the Fast Fourier Transform.

## References

- W.M.P. van der Aalst. *Timed coloured Petri nets and their application to logistics*. PhD thesis, Eindhoven University of Technology, Eindhoven, 1992.
- W.M.P. van der Aalst. Interval Timed Coloured Petri Nets and their Analysis. In M. Ajmone Marsan, editor, *Application and Theory of Petri Nets 1993*, volume 691 of *Lecture Notes in Computer Science*, pages 453-472. Springer-Verlag, Berlin, 1993.
- W.M.P. van der Aalst. Using Interval Timed Coloured Petri Nets to Calculate Performance Bounds. In G. Haring and G. Kotsis, editors, *Proceedings of the 7th International Conference of Modelling Techniques and Tools for Computer Performance Evaluation*, volume 794 of *Lecture Notes in Computer Science*, pages 425-444. Springer-Verlag, Berlin, 1994.
- W.M.P. van der Aalst. The Application of Petri Nets to Workflow Management. *The Journal of Circuits, Systems and Computers*, 8(1):21-66, 1998.
- W.M.P. van der Aalst. Workflow Verification: Finding Control-Flow Errors using Petri-net-based Techniques. In *Business Process Management: Models, Techniques, and Empirical Studies*, Springer-Verlag, Berlin, 2000.
- W.M.P. van der Aalst and M.A. Odijk. Analysis of Railway Stations by means of Interval Timed Coloured Petri Nets. *Real-Time Systems*, 9(3):241-263, 1995.
- G. Balbo and M. Silva, editors. *Performance Models for Discrete Event Systems with Synchronisations: Formalisms and Analysis Techniques*, Zaragoza, Sept 1998. Kronos.
- J.A. Bergstra and J.W. Klop. Process Algebra for Synchronous Communication. *Information and Control*, 60(1-3):109-137, 1984.
- B. Berthomieu and M. Diaz. Modelling and verification of time dependent systems using Time Petri Nets. *IEEE Transactions on Software Engineering*, 17(3):259-273, March 1991.
- B. Berthomieu and M. Menasche. An enumerative approach for analyzing time Petri nets. In R.E.A. Mason, editor, *Information Processing: proceedings of the IFIP congress 1983*, volume 9 of *IFIP congress series*, pages 41-46. Elsevier Science Publishers, Amsterdam, 1983.
- J. Carlier and P. Chretienne. Timed Petri net schedules. In G. Rozenberg, editor, *Advances in Petri Nets 1988*, volume 340 of *Lecture Notes in Computer Science*, pages 62-84. Springer-Verlag, Berlin, 1988.
- G. Chiola, C. Dutheillet, G. Franceschinis, and S. Haddad. On well-formed coloured nets and their symbolic reachability graph. In *Proceedings of the 11th International Conference on Applications and Theory of Petri Nets*, pages 387-411, Paris, June 1990.
- P. Chretienne. *Les réseaux de petri temporisés*. PhD thesis, Univ. Paris VI, Paris, 1983.
- C. Dutheillet and S. Haddad. Regular Stochastic Petri Nets. In *Proceedings of the 10th International Conference on Applications and Theory of Petri Nets*, Bonn, June 1989.
- C.A. Ellis and G.J. Nutt. Modelling and Enactment of Workflow Systems. In M. Ajmone Marsan, editor, *Application and Theory of Petri Nets 1993*, volume 691 of *Lecture Notes in Computer Science*, pages 1-16. Springer-Verlag, Berlin, 1993.
- G. Florin and S. Natkin. Evaluation based upon Stochastic Petri Nets of the Maximum Throughput of a Full Duplex Protocol. In C. Girault and W. Reisig, editors, *Application and*

- theory of Petri nets: selected papers from the first and the second European workshop*, volume 52 of *Informatik Fachberichte*, pages 280-288, Berlin, 1982. Springer-Verlag, Berlin.
- C. Ghezzi, D. Mandrioli, S. Morasca, and M. Pezze. A unified high-level Petri net formalism for time-critical systems. *IEEE Transactions on Software Engineering*, 17(2):160-172, Feb 1991.
- K.M. van Hee. *Information System Engineering: a Formal Approach*. Cambridge University Press, 1994.
- K.M. van Hee, L.J. Somers, and M. Voorhoeve. Executable specifications for distributed information systems. In E.D. Falkenberg and P. Lindgreen, editors, *Proceedings of the IFIP TC 8 / WG 8.1 Working Conference on Information System Concepts: An In-depth Analysis*, pages 139-156, Namur, Belgium, 1989. Elsevier Science Publishers, Amsterdam.
- M.A. Holliday and M.K. Vernon. A Generalised Timed Petri Net Model for Performance Analysis. *IEEE Transactions on Software Engineering*, 13(12):1279-1310, Dec 1987.
- S. Jablonski and C. Bussler. *Workflow Management: Modeling Concepts, Architecture, and Implementation*. International Thomson Computer Press, London, UK, 1996.
- K. Jensen. *Coloured Petri Nets. Basic Concepts, Analysis Methods and Practical Use*. EATCS monographs on Theoretical Computer Science. Springer-Verlag, Berlin, 1996.
- P. Lawrence, editor. *Workflow Handbook 1997, Workflow Management Coalition*. John Wiley and Sons, New York, 1997.
- C. Lin and D.C. Marinescu. On Stochastic High-Level Petri Nets. In *Proceedings of the International Workshop on Petri Nets and Performance Models*, pages 34-43. IEEE Computer Society Press, Madison, 1987.
- M. Ajmone Marsan. Stochastic Petri Nets: An Elementary Introduction. In G. Rozenberg, editor, *Advances in Petri Nets 1989*, volume 424 of *Lecture Notes in Computer Science*, pages 1-29. Springer-Verlag, Berlin, 1990.
- M. Ajmone Marsan, G. Balbo, A. Bobbio, G. Chiola, G. Conte, and A. Cumani. On Petri Nets with Stochastic Timing. In *Proceedings of the International Workshop on Timed Petri Nets*, pages 80-87, Torino, 1985. IEEE Computer Society Press.
- M. Ajmone Marsan, G. Balbo, and G. Conte. A Class of Generalised Stochastic Petri Nets for the Performance Evaluation of Multiprocessor Systems. *ACM Transactions on Computer Systems*, 2(2):93-122, May 1984.
- M. Ajmone Marsan, G. Balbo, and G. Conte. *Performance Models of Multiprocessor Systems*. The MIT Press, Cambridge, 1986.
- M. Ajmone Marsan, G. Balbo, and G. Conte et al. *Modelling with Generalized Stochastic Petri Nets*. Wiley series in parallel computing. Wiley, New York, 1995.
- M. Ajmone Marsan and G. Chiola. On Petri Nets with Deterministic and Exponentially Distributed Firing Times. In G. Rozenberg, editor, *Advances in Petri Nets 1987*, volume 266 of *Lecture Notes in Computer Science*, pages 132-145. Springer-Verlag, Berlin, 1987.
- P. Merlin. *A Study of the Recoverability of Computer Systems*. PhD thesis, University of California, Irvine, California, 1974.
- P. Merlin and D.J. Faber. Recoverability of communication protocols. *IEEE Transactions on Communications*, 24(9):1036-1043, Sept 1976.
- M.K. Molloy. *On the Integration of Delay and Throughput Measures in Distributed Processing Models*. PhD thesis, University of California, Los Angeles, 1981.



- S. Morasca, M. Pezzè, and M. Trubian. Timed High-Level Nets. *The Journal of Real-Time Systems*, 3:165-189, 1991.
- T. Murata. Petri Nets: Properties, Analysis and Applications. *Proceedings of the IEEE*, 77(4):541-580, April 1989.
- C.A. Petri. *Kommunikation mit Automaten*. PhD thesis, Institut für instrumentelle Mathematik, Bonn, 1962.
- C.V. Ramamoorthy and G.S. Ho. Performance Evaluation of Asynchronous Concurrent Systems Using Petri Nets. *IEEE Transactions on Software Engineering*, 6(5):440-449, Sept 1980.
- C. Ramchandani. *Performance Evaluation of Asynchronous Concurrent Systems by Timed Petri Nets*. PhD thesis, Massachusetts Institute of Technology, Cambridge, 1973.
- R.R Razouk and C.V. Phelps. Performance analysis using Timed Petri Nets. In *Proceedings of the 1984 International Conference on Parallel Processing*, pages 126-128. IEEE Computer Society Press, Ohio, 1984.
- W. Reisig and G. Rozenberg, editors. *Lectures on Petri Nets I: Basic Models*, volume 1491 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 1998.
- W. Reisig and G. Rozenberg, editors. *Lectures on Petri Nets II: Applications*, volume 1492 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 1998.
- J. Sifakis. Use of Petri Nets for performance evaluation. In H. Beilner and E. Gelenbe, editors, *Proceedings of the Third International Symposium IFIP W.G. 7.3., Measuring, modelling and evaluating computer systems (Bonn-Bad Godesberg, 1977)*, pages 75-93. Elsevier Science Publishers, Amsterdam, 1977.
- J. Sifakis. Performance Evaluation of Systems using Nets. In W. Brauer, editor, *Net theory and applications : Proceedings of the advanced course on general net theory, processes and systems (Hamburg, 1979)*, volume 84 of *Lecture Notes in Computer Science*, pages 307-319. Springer-Verlag, Berlin, 1980.
- R. Valette. Analysis of Petri Nets by Stepwise Refinements. *Journal of Computer and System Sciences*, 18:35-46, 1979.
- C.Y. Wong, T.S. Dillon, and K.E. Forward. Timed Places Petri Nets with Stochastic Representation of Place Time. In *Proceedings of the International Workshop on Timed Petri Nets*, pages 96-103, Torino, 1985. IEEE Computer Society Press.
- A. Zenie. Coloured Stochastic Petri Nets. In *Proceedings of the International Workshop on Timed Petri Nets*, pages 262-271, Torino, 1985. IEEE Computer Society Press.
- W.M. Zuberek. Timed Petri Nets and Preliminary Performance Evaluation. In *Proceedings of the 7th annual Symposium on Computer Architecture*, volume 8(3) of *Quarterly Publication of ACM Special Interest Group on Computer Architecture*, pages 62-82, 1980.