

March 1985

LIDS-P-1443

PERFORMANCE EVALUATION OF DECISION-FREE
STOCHASTIC TIMED PETRI NETS

by

R. Paul Wiley
Robert R. Tenney

Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, MA 02139

This research was supported by the Office of Naval Research under grants
ONR/N00014-77-C-0532 (NR 041-519) and ONR/N00014-84-K-0519 (NR 649-003).

INTRODUCTION

The advent of low-cost processors has made the construction of highly complex, decentralized systems feasible. These systems are usually divided into interacting components which operate concurrently and coordinate their operations by using an asynchronous protocol. Petri Nets have frequently been used to model such systems since they can model concurrency and asynchronous protocols [11]. Specific examples of systems modeled by Petri Nets include computer systems [12], communication protocols [5] and manufacturing systems [6].

One aspect of designing these complex, concurrent, asynchronous systems is predicting their performance with respect to time related measures. These measures could include, but are not limited to, throughput, backlog, the probability that a shared resource is free at any given time, the average time it takes to complete each particular task, etc. In order to use Petri Nets to predict these performance measures, the notion of time must be added to the original definition. Several authors, Ramchandani [13], Sifakis [14], and Ramamoorthy and Ho [12], have studied Deterministic Timed Petri Nets (DTPNs), where deterministic processing times have been added to model fixed delays. Stochastic Timed Petri Nets (STPNs) have been studied by Zuberek [15], Molloy [9], and Marsan, et. al. [8]. All of these authors perform the analysis by transforming the problem into an equivalent Markov Chain, whose steady state probability distribution is then found by conventional techniques. Since the numbers of states in the chain tends to grow exponentially with the size of the net, their techniques are limited to solving relatively small problems.

In this paper, we will show how the performance of a simple subclass of STPNs, that of Decision-Free STPNs (DFSTPNs), can be predicted. This subclass can model concurrency and coordination, but not decisions, and have been used to model certain manufacturing systems [6]. The major contribution is a new analysis technique which uses a set of equations that describe the behavior of the system. Cohen, et.al [3] and Dubois and Stecke [6] use similar equations to analyze the deterministic case. The resulting algorithms are computationally more attractive and the general approach can be extended to study less restrictive classes of STPNs.

An overview of this paper goes as follows. We start by defining STPNs and presenting the relevant concepts. Then, the study is limited to DFSTPNs. We examine three related processes which describe firing times, relative firing times, and intertransition times. The first process introduces the new analytical technique, the second one provides a convergence proof, and the third one allows performance measures of interest to be computed.

II. STOCHASTIC TIMED PETRI NETS

Stochastic Timed Petri Nets are graphs with two types of nodes: places (drawn as a circle and labelled p_i) and transitions (drawn as a line segment and labelled t_i), and directed arcs going from one type of node to the other (see Figure 1). Besides nodes and arcs, a STPN assigns, to every place, a nonnegative number of tokens, called the marking of the net, and a nonnegative random processing time¹. When a token arrives at a place, it is defined to be unavailable (following Sifakis [14]). The token remains in this state until the instant when the processing has finished. At this time, the token becomes available. Initially, all tokens are assumed to be unavailable.

Tokens move around by transition firings. A transition is enabled, that is, it may fire, only when all of its input places have available tokens. When a transition does fire, it removes a token from all of its input places and adds a token to all of its output places. Two or more transitions are said to be in conflict if firing one will disable the others. If an enabled transition is not in conflict, it fires instantly. If several enabled transitions are in conflict, then a decision rule, specified a priori, selects one and that transition fires instantly. Thus, once the graph, the initial marking, the processing times and the decision rules are specified, the STPN evolves autonomously in time.

One can think of the operation of a STPN as a succession of markings, with the transitions from one marking to another specified by the transition

¹Many authors, Ramchandani [13], Zuberek [15], and Molloy [9], associate processing times with transitions instead of places. The distinction makes no difference for the class of problems studied in this paper.

firings and the rules for moving tokens. If the number of tokens in any particular place can never exceed one, regardless of what processing time probability distributions and decision rules are assigned, the STPN is said to be safe. If there exists a τ_j for any transition t_j and for any $\epsilon > 0$ such that

$$\text{Prob}\{\text{transition } t_j \text{ will fire before } \tau_j\} > 1 - \epsilon$$

regardless of what marking the net has reached, the STPN is defined to be live. And finally, if there exists a directed path from any node (place or transition) to any other node, the STPN is strongly connected. Throughout this paper, we will assume that the STPNs we are studying are strongly connected, live, safe, and with a finite number of nodes.

III. DECISION-FREE STOCHASTIC TIMED PETRI NETS

Decision-Free Stochastic Timed Petri Nets are STPNs in which every place has exactly one input and one output transition (see Figure 2 for an example). Thus, transitions are never in conflict and decisions rules are not needed. This fact makes DFSTPNs much easier to analyze than broader subclasses of STPNs, but also imposes a modeling limitation since, as noted previously, they cannot model decisions. They can, however, model concurrency and coordination. The performance measures that we are interested in obtaining are the steady state firing rate of every transition, the steady state number of tokens in every place, and the steady state delay a token encounters in every place.

But before we proceed, let us review the pertinent results of Decision-Free Petri Net (DFPN) theory, or DFSTPNs without the processing times. These Petri Nets, also called Marked Graphs in the literature, have been extensively studied [3], [9], with most of the work centering around establishing liveness and safeness conditions. In the following list of results, directed circuits are directed paths from a node back to itself. The results that are important for our study are:

- (i) The number of tokens in a directed circuit does not change as a result of transition firings.
- (ii) A DFPN is live iff the number of tokens in every directed circuit is positive.

- (iii) A live marking is safe iff every place is contained in a directed circuit with a token count of one.

We will now show that these results also hold true for DFSTPNs by showing that they are live and safe iff the underlying DFPNs are live and safe. The safeness equivalence is true by definition. The liveness equivalence, however, must be proved. To that effect let $z_i(k)$ be the processing time for place p_i for the k^{th} token to arrive at that place. We will assume that there exists a τ such that

$$\text{Prob}\{z_i(k) < \tau\} > 1-\varepsilon$$

for any $\varepsilon > 0$, regardless of the numerical values of the other service times (a mild independence condition that is satisfied, for example, by non-negative, bounded random variables).

Proposition 1. Given the processing time assumption mentioned above, a DFSTPN is live iff the underlying DFPN is live.

Proof: If the DFSTPN is live, then there exists at least one sequence of transition firings and a τ such that, for any transition t_j and any $\varepsilon > 0$,

$$\text{Prob}\{\text{transition } t_j \text{ will fire before } \tau\} > \varepsilon.$$

The same sequence of transition firings is valid in the underlying DFPN, which shows it is live.

Assume the underlying DFPN is live. The liveness proof given by

Commoner, et.al. [4] ensures that, for any particular transition t_j , there exists a circuit-free backtracking subgraph such that at least one of the transitions, call it t_i , contains tokens in all of its input places. By the assumption on the processing times, there exists a τ_i such that, for any $\varepsilon > 0$,

$$\text{Prob}\{\text{transition } t_i \text{ will fire before } \tau_i\} > 1-\varepsilon.$$

A similar argument can then be given for all other transitions along the subgraph. Thus, there exists a τ_j for any $\varepsilon > 0$ such that

$$\text{Prob}\{\text{transition } t_j \text{ will fire before } \tau_j\} > 1-\varepsilon$$

which shows the DFSTPN is live.

Q.E.D.

As can be noted from the listed results, directed circuits play an important role in the study of DFSTPNs. It is useful to think of the operation of a DFSTPN in terms of an equivalent circuit-free unfolded DFSTPN. This unfolded net consists of an infinite number of identical stages. Each stage is formed by breaking the regular DFSTPN immediately before every place containing a token in the initial marking, and then "unfolding" the net. For example, Figure 3 shows the Unfolded DFSTPN that corresponds to the DFSTPN shown in Figure 2.

A formal procedure to construct one stage of the Unfolded DFSTPN is given by Algorithm 1. In this algorithm, the variable k represents the

stage and the transition firings are symbolic, not actual transition firings of the DFSTPN. We also denote the set of input (output) places to transition t_i by $\cdot t_i$ ($t_i \cdot$), as commonly done in the literature [13], [14].

Algorithm 1: Construction of one stage of the Unfolded DFSTPN

INITIALIZED: M to M_0 (initial marking)

- $T_M^k \leftarrow \{t_i \mid t_i \text{ is enabled by } M\}$
- $P_k \leftarrow \{p_j \mid p_j \text{ contains a token in } M\}$
- Create a place p_j^k for all $p_j \in P_k$

DO UNTIL $T_M^k = \emptyset$

- For every transition $t_i \in T_M^k$
 - (a) Create a transition t_i^k
 - (b) Direct an arc from p_j^k to t_i^k for every $p_j \in \cdot t_i$
 - (c) Create a place for every $p_j \in t_i \cdot$

IF $p_j \in P_k$

THEN

Label the place p_j^{k+1}

ELSE

Label the place p_j^k

$P_k \leftarrow P_k \cup p_j$

ENDIF

- Fire all transitions $t \in T_M^k$

- Update M , T_M^k

ENDUNTIL

It can be shown by using an argument similar to the liveness proof given by Commoner, et.al [4] that Algorithm 1 fires every transition exactly once. We also note that the marking at the end of the firing sequence is the same as the initial marking since every transition has been fired the same number of times (Murata [10], Property 2). Thus, all stages are identical except for the variable k .

Besides constructing a stage of the Unfolded DFSTPN, Algorithm 1 also provides a partial order among transition firings. Before a transition can fire, it must have tokens in all of its input places. When a transition comes up in the partial order established by the algorithm, all of its input places do, in fact, contain tokens.

Visualizing the operation of DFSTPN in terms of its unfolded equivalent is not only convenient, but it is also exact. We will show that the two systems are dynamically equivalent after we derive the firing time state equations, which we do next.

IV. FIRING TIME STATE EQUATIONS

Given a DFSTPN, a natural question to ask is the time at which the k th firing of each transition will occur. In this section, we derive equations to answer this question. Similar equations are derived in Cohen, et.al. [3] to analyze bounded² Decision-Free DTPNs. Their derivation is not applicable to DFPNs with random processing times. Dubois and Stecke [6] also use similar equations, but do not derive them. Since these equations are important for this paper, we present a brief derivation and show that they are consistent.

To this objective, consider the DFSTPN shown in Figure 4 and assume that the initial marking does not place a token in either p_1 or p_2 . Then, if we let $x_i(k)$ be the time of the k^{th} firing of transition t_i ,

$$x_3(k) = \max\{z_1(k) + x_1(k), z_2(k) + x_2(k)\} \quad (1)$$

where $z_i(k)$ is the processing time of place p_i for the k^{th} token that arrives to that place. Note that this type of equation can be written for a DFSTPN even if it is not live or strongly connected.

Now, assume that the net shown in Figure 4 is part of a larger DFSTPN that is strongly connected, live and safe. Equation (1) holds as long as the initial marking does not place a token in p_1 or p_2 . If the initial marking places a token in p_1 , but not p_2 , then

²In the stochastic case, tokens are free to cross each other in a place depending on the numerical value of their random processing time. To avoid keeping track of all possible orderings of tokens, we must assume safeness.

$$x_3(k) = \max\{z_1(k) + x_1(k-1), z_2(k) + x_2(k)\} \quad (2)$$

where $x_1(0) = 0$. This equation is consistent since the liveness and safeness assumptions ensure that p_1 is contained in a directed circuit containing one token, which implies that $x_1(k-1) \leq x_3(k) \leq x_1(k)$ for all k . Similarly, if the initial marking places tokens in both p_1 and p_2 , then we conclude that

$$x_3(k) = \max\{z_1(k) + x_1(k-1), z_2(k) + x_2(k-1)\} \quad (3)$$

In this discussion, we have considered a transition with only two input places. If there had been more, than the maximization operator of the state equation would have included more terms, but could still be easily written. Thus, firing time state equations could be obtained for every transition in a DFSTPN in the manner discussed above. Figure 5 gives an example of a DFSTPN and its corresponding firing time state equations.

Given the processing times, the state equations could be used to calculate the transition firing times recursively. These calculations would have to be performed in the proper order to satisfy all the causality constraints. As noted previously, any order consistent with the partial order provided by Algorithm 1 satisfies these requirements. The calculations would also have to be initialized properly. This can be done by setting $x_i(0) = 0$ for all places p_i . This is consistent with our assumption that, initially, all tokens are unavailable.

We mentioned previously that the visualization of the operation of a

DFSTPN in terms of its equivalent Unfolded DFSTPN was not only convenient, but also exact. We now elaborate. By using the state equations, we will show that the two systems are dynamically equivalent. That is, we will show that there exists a one to one correspondence between the transition firings and processing times of the two systems such that the related transition firings occur at the same time provided the related processing times are equal. To this effect, let G be a DFSTPN and UG be its corresponding Unfolded DFSTPN. Then, we specify that the initial marking in UG is a token in every place p_i^1 , where p_i contains a token in the initial marking of G . Now for the proposition.

Proposition 2: The dynamic evolution of a DFSTPN is equivalent to its corresponding Unfolded DFSTPN.

Proof: If we associate $x_j(k)$ and $z_j(k)$ with the transition t_j^k and place p_j^k of the Unfolded DFSTPN, respectively, then the state equations for the DFSTPN and its corresponding Unfolded DFSTPN are identical.

Q.E.D.

The dynamic behavior of a DFSTPN is difficult to understand based on the state equations as given. To improve this situation, we will show that the state equations can be expanded so that the k^{th} firing time of any transition depends only on the $k-1$ firing times of other transitions. Then we will show that the resulting equations can be written in matrix form by converting them to a $(\max,+)$ algebra.

As seen from the derivation of the firing time state equations, the k^{th}

firing of any particular transition t_j depends, in general, on the $k-1$ or k^{th} firings of other transitions. By repeatedly substituting the k^{th} firings times of these other transitions, the equations can be expanded so that they depend only on the $k-1$ firing times of other transitions and on the k^{th} service times of several places. The repeated substitution procedure corresponds to tracing backwards along all paths in the Unfolded DFSTPN, starting from t_j^k and ending at transitions in the previous stage. Since the Unfolded DFSTPN contains no directed circuits, these backtracking paths are guaranteed to end in the previous stage and are finite. And all places that appear along the path correspond to the k^{th} stage of the Unfolded DFSTPN. Therefore only $z_i(k)$ terms appear in the state equations for $x_j(k)$, and no $z_i(m)$, $m < k$, terms appear.

To make the structure of these equations clearer, substitute the max operation by \oplus and + by \odot , as done in [2], [3], and [6]. If we let $\underline{x}(k)$ be the vector for the k^{th} firing time of all transitions, then the discussion on repeated substitution and some algebraic manipulations shows that

$$\underline{x}(k) = A[\underline{z}(k)] \odot \underline{x}(k-1) \quad (4)$$

where $A[\underline{z}(k)]$ is a matrix whose components are functions of $\underline{z}(k)$, the vector of processing time variables, and the matrix multiplication is in the sense of \oplus and \odot operators. The assumption that all tokens are initially unavailable implies that $\underline{x}(0) = \underline{0}$, the vector of all zeroes.

Cohen, et.al. [2] analyzes systems which have state equations of the form given in (4), but with deterministic processing times. In Cohen,

et.al. [3] this theory is applied to Decision-Free DTPNs. These papers show that the firing times of all transitions will become periodic modulo a constant, regardless of what numerical value the initial firing time had. That is, there exists a finite k_0 , λ and d such that for any $\underline{x}(0)$

$$\underline{x}(k+d) = d\lambda\underline{1} + \underline{x}(k) \quad (5)$$

for all $k \geq k_0$. In this previous equation, $\underline{1}$ is a vector of ones and λ corresponds to the maximal average circuit delay or the maximum, over all directed circuits, sum of delays divided by the number of tokens, of each circuit³. The directed circuits with average circuit delay of λ are denoted as critical circuits⁴. If the critical circuit is unique and contains only one token, then d in equation (5) is equal to 1 and furthermore, the vector $\underline{x}(k)$ which is attained for k sufficiently large is unique to within an additive constant. We shall make use of this result later.

So far, in this section, we have spent a lot of time obtaining and clarifying the structure of the firing time state equations. Since the processing times are, in general, random variables, it is necessary that we try to characterize the probability distribution for $\underline{x}(k)$, which we denote by $f_{\underline{x}(k)}(\underline{X})$. To keep the calculations from getting arbitrarily complex, we assume, from this point on, that the processing times are independent random

³The mentioned references proved their results on conditions of $G(A)$, the graph whose weighted incidence matrix is A . In this paper, we translate those conditions to conditions on the Decision-Free DTPN being considered.

⁴Thus, the long term behavior of a Decision -Free DTPN is determined by one or more of these critical circuits. This fact alludes to the major difficulty in studying the stochastic case: the "critical" circuits are, in a sense, constantly changing in time.

variables over the stages. With this assumption, we can show the following.

Proposition 3. If the $\underline{z}(k)$ are independent random vectors over k , and independent from $\underline{x}(0)$, then the process $\{\underline{x}(k), k \in \mathbb{N}\}$ is Markov.

Proof: Follows from equation (4) and the independence assumptions.

Q.E.D.

Thus, the vector of firing times, $\underline{x}(k)$, constitutes the state of a Markov Process. This conclusion is intuitive if one realizes that the transitions whose firing times are included in $\underline{x}(k)$ constitute a cut set in the Unfolded DFSTPN. The Markov property also implies that one can use the firing time state equation to calculate $f_{\underline{x}(k)}(\underline{X})$ starting from $f_{\underline{x}(0)}(\underline{X})$ for all k . This, however, may not be the indicated course of action. The variance for any particular $x_i(k)$ will increase without bound as k increases, as we can convince ourselves by considering a very simple DFSTPN containing one transition and one place connected in a loop with a random processing time. Furthermore, we know of no way in which to characterize $f_{\underline{x}(k)}(\underline{X})$ as k goes to infinity by using some central limit theorem or assuming some clever probability for the processing times.

Thus, we cannot use $f_{\underline{x}(k)}(\underline{X})$ to obtain the desired performance measures since the probability distributions do not converge. To correct this situation, we will do a few things. First, we will derive the relative firing times process and show that the probability distributions that characterize it will converge by making some additional assumptions on the

processing times. Then, we will define the intertransition times process, show that the probability distributions that characterize it will converge by relating it to the relative firing times. And finally, we will show how the performance measures can be obtained from the steady state probability distributions of the intertransition times.

V. RELATIVE FIRING TIMES

The relative firing times are obtained by performing a transformation on the firing times. To that effect, pick a transition t_r to be the reference transition. Then $\underline{w}(k)$, the vector of relative firing times, is defined by

$$\underline{w}(k) = \underline{x}(k) - x_r(k)\underline{1} \quad (6)$$

where $\underline{1}$ is a vector of ones and $x_r(k)$ is the k^{th} firing time of the reference transition. We can also represent this transformation in $(\max, +)$ algebra as follows

$$\underline{w}(k) = [x_r(k)]^{-1} \odot \underline{x}(k) \quad (7)$$

where $[\cdot]^{-1}$ denotes subtraction in the usual algebra. As can be seen from this last equation, the transformation performed on $\underline{x}(k)$ to obtain $\underline{w}(k)$ is a scaling operation in the $(\max, +)$ algebra domain. In the paragraphs that follow, to clarify the discussion, we let

$$\underline{w}(k) = S[\underline{x}(k)] \quad (8)$$

represent the transformation indicated by equations (6) and (7).

In order for this transformation from $\underline{x}(k)$ to $\underline{w}(k)$ to be useful, it must be well defined. That is, we must show that the derivation of $\underline{w}(k+1)$ from $\underline{x}(k+1)$ using the definition is equivalent to deriving $\underline{w}(k+1)$ from $\underline{w}(k)$ using a state equation for the relative firing times process. We do this

next.

Proposition 4. The transformation from $\underline{x}(k)$ to $\underline{y}(k)$ is well defined, that is,

$$S[\underline{x}(k+1)] = S\{A[\underline{z}(k)] \odot S[\underline{x}(k)]\} \quad (9)$$

Proof⁵: By the definition of relative firing times, we have that

$$\begin{aligned} S[\underline{x}(k+1)] &= A[\underline{z}(k)] \odot \underline{x}(k) \odot \{ \underline{a}_r^T[\underline{z}(k)] \odot \underline{x}(k) \}^{-1} \odot \underline{x}_r(k) \odot \underline{x}_r^{-1}(k) \\ &= A[\underline{z}(k)] \odot \underline{x}(k) \odot \underline{x}_r^{-1}(k) \odot \{ \underline{a}_r^T[\underline{z}(k)] \odot \underline{x}(k) \odot \underline{x}_r^{-1}(k) \}^{-1} \\ &= A[\underline{z}(k)] \odot S[\underline{x}(k)] \odot \{ \underline{a}_r^T \underline{z}(k) \odot S[\underline{x}(k)] \}^{-1} \\ &= S\{A[\underline{z}(k)] \odot S[\underline{x}(k)]\} \end{aligned} \quad (10)$$

which is what we wanted to show.

Q.E.D.

Expanding the last equation, we obtain the state equation for the relative firing times process

⁵This proof uses (max,+) algebra results which can be verified by converting back to the usual algebra or by consulting [2].

$$\underline{w}(k+1) = A[\underline{z}(k)] \odot \underline{w}(k) \odot \{ \underline{a}_r^T[\underline{z}(k)] \odot \underline{w}(k) \}^{-1} \quad (11)$$

With this state equation, we can prove the following proposition.

Proposition 5. If the $\underline{z}(k)$ are independent random vectors over k , and independent from $\underline{w}(0)$, then the process $\{\underline{w}(k), k \in \mathbb{N}\}$ is Markov.

This conclusion is not surprising if we remember that the transitions whose relative firing times are contained in $\underline{w}(k)$ form a cut set of the Unfolded DFSTPN. The transformation $S(\cdot)$ subtracts the absolute value of time, and this has no effect on the Markov property. The transformation does, however, provide us with some nice convergence properties, as we will see next.

VI. PROPERTIES OF THE RELATIVE FIRING TIMES MARKOV PROCESS

In this section, we analyze the consequences of making further assumptions on the processing times. First, if we assume that the processing times are bounded, we will show that the state space is bounded. In turn, if we assume the processing time probability distributions are stationary and in discrete time, we will show that the resulting Markov process has a single recurrence class. Then, we will combine all of the assumptions stated above, and analyze the consequences.

So, first assume that the service times are bounded for all places and for all k . By incorporating this assumption into the liveness proof given by Commoner, et.al. [4], we can show that the time to fire any particular transition, from all possible markings, is bounded. Call this bound τ . We can also show that a strongly connected DFSTPN can be considered as the union of a finite number of directed circuits, each containing one token. This, in turn, leads us to conclude that the number of times any particular transition fires differs from the number of times any other transition fires by less than some finite number. Let this number be L . Now, consider the DFSTPN at any instant of time when the reference transition t_r fires, say for the k^{th} time. Let t_j be any other transition. We know that the number of times transition t_j has fired differs from k by at most L . We also know that the time between t_j firings is at most τ . Therefore

$$|x_r(k) - x_j(k)| \leq \tau L \quad (12)$$

Thus, every element of $\underline{w}(k)$ is bounded. We have just shown the following:

Proposition 6. If all elements of $\underline{z}(k)$ are bounded for all k , then all elements of $\underline{w}(k)$ are bounded for all k .

Now assume that the processing times $\underline{z}(k)$ are identically distributed over k and in discrete time. Further assume that there exists a value of \underline{Z} for which $f_{\underline{z}}(\underline{Z}) > 0$ and for which the DFSTPN, with that particular value of \underline{Z} , has a unique critical circuit containing one token. If that value of \underline{Z} occurs enough times in a row, then the results of Cohen, et. al. cited in Section IV tell us that there exists a finite k_0 such that for any $x(0)$

$$\underline{x}(k+1) = \lambda \underline{1} + \underline{x}(k) \quad (13)$$

for $k \geq k_0$ and that this value of $\underline{x}(k)$ is unique to within an additive constant. The values of $\underline{x}(k)$ related by equation (13) are, in fact, the same state in the relative firing times Markov process. Thus, if this value of \underline{Z} occurs k_0 times in a row, an event which occurs with probability $[f_{\underline{z}}(\underline{Z})]^{k_0} > 0$, the relative firing times process is in a recurrent state and will remain in this state as long as the value of \underline{Z} keeps occurring. Since this state is reachable from anywhere in the state space, we have just shown the following:

Proposition 7: If $\underline{z}(k)$ is independent and identically distributed over k and in discrete time, and if there exists a value of \underline{Z} for which $f_{\underline{z}}(\underline{Z}) > 0$ such that the DFSTPN has a unique critical circuit containing only one token for that particular value of \underline{Z} , then the discrete-state discrete-index Markov process $\{\underline{w}(k), k \in \mathbb{N}\}$ has a single recurrence class and a recurrent

state with a self-loop.

In this subsection, we have shown that the $\underline{w}(k)$ state space is bounded if the $\underline{z}(k)$ are bounded for all k . If we also assume that the $\underline{z}(k)$ are in discrete time, then we conclude that the number of $\underline{w}(k)$ states is finite. An finally, if we also assume that the $\underline{z}(k)$ are stationary and that there exists a value of \underline{Z} for which $f_{\underline{z}}(\underline{Z}) > 0$ and for which the DFSTPN has a unique critical circuit containing only one token, we conclude that the resulting $\{\underline{w}(k), k \in \mathbb{N}\}$ discrete-state discrete-index Markov process has a single recurrence class, a finite number of states, and at least one recurrent state with a self-loop. Such Markov chains have a unique ergodic probability distribution [1], which we will denote by $F(\underline{W})$.

Thus, we have shown that the relative firing times Markov process has an ergodic probability distribution if we make the given assumptions. The problem remains, however, of how to calculate the performance measures. We will obtain these measures from the intertransition times, which we define next. We will show that the intertransition times also have an ergodic probability distribution by relating them to the relative firing times.

VII. INTERTRANSITION TIMES STATE EQUATIONS

In intertransition time is the time between the firings of two transitions that have one place in between them, that is, it is the time a token spends in a place. Consider the DFSTPN shown in Figure 4, which we assume is part of a larger, live and safe, DFSTPN. If we assume that the initial marking does not place a token in p_1 , then the time between the k th firings of transitions t_1 and t_3 is given by

$$\begin{aligned} v_{13}(k) &= x_3(k) - x_1(k) \\ &= \max\{z_1(k), z_2(k) + x_2(k) - x_3(k-1) + x_3(k-1) - x_1(k)\} \end{aligned} \quad (14)$$

where we have used the firing time state equations and added a zero to one of the terms in the maximization operator. Now, place p_2 is contained in a directed circuit containing one token. By using this fact, we can conclude that $x_2(k) - x_3(k-1)$ can be expressed as a sum of intertransition times along the specified directed circuit. Similarly, $x_1(k) - x_3(k-1)$ can be expressed as a sum of intertransition times by noting that place p_1 is contained in a directed circuit with one token. Thus, we conclude that $v_{13}(k)$ can be expressed in terms of the processing times corresponding to the input places of a transition t_3 , and in terms of the intertransition times along circuits with one token containing transition t_3 .

The reasoning used above can be used to write intertransition time state equations among any pair of transitions. More input places can be handled by using the same argument given above several times. And if place p_1 contains a token in the initial marking, then the definition of the

intertransition changes slightly to become

$$v_{13}(k) = x_3(k) - x_1(k-1) \quad (15)$$

The rest of the reasoning is identical.

As an example, consider the DFSTPN shown in Figure 5. The state equations for the four possible intertransition times are

$$v_{11}(k) = \max\{z_1(k), z_3(k) + v_{12}(k-1)\} \quad (16)$$

$$v_{21}(k) = \max\{z_3(k), z_1(k) - v_{12}(k-1)\} \quad (17)$$

$$v_{12}(k) = \max\{z_2(k), z_4(k) - v_{21}(k)\} \quad (18)$$

$$v_{22}(k) = \max\{z_4(k), z_2(k) + v_{21}(k)\} \quad (19)$$

As with the firing time state equations, the intertransition times can be calculated recursively by doing the calculations in the proper order and by initializing the process properly. The proper order can be achieved by using the partial order established by Algorithm 1. When a transition t_j is fired by that algorithm, all the information required to calculate $v_{ij}(\cdot)$ is known for all possible t_i . To initialize the intertransition time state equations, one must refer back to the original definitions and remember that $\underline{x}(0) = \underline{0}$.

Since the intertransition times can be calculated recursively, we can conclude that there exists a transformation $g(\cdot)$ such that

$$\underline{v}(k) = g\{\underline{z}(k), \underline{v}(k-1)\} \quad (20)$$

In this last equation, $\underline{v}(k)$ is the vector of intertransition times where $[\underline{v}(k)]_j = v_{ij}(k)$ if place p_j is between transitions t_i and t_j . The exact nature of $g(\cdot)$ is very complicated, but the existence of such a transformation is important since it can be used to show the following:

Proposition 8. If the $\underline{z}(k)$ are independent random vectors over k , and independent from $\underline{v}(0)$, then the $\{\underline{v}(k), k \in \mathbb{N}\}$ process is Markov.

We derived the intertransition times since the firing times process did not converge and since we could not obtain the performance measures directly from the relative firing times. Next we will show that the state probability distributions for the intertransition times will converge by relating them to the relative firing times, and then we will show how the performance measures can be calculated.

VIII. TRANSFORMATION FROM RELATIVE TO INTERTRANSITION FIRING TIMES

From the previous section, we note that the intertransition time vector, $\underline{v}(k)$, is defined in terms of $\underline{x}(k)$ and $\underline{x}(k-1)$. This implies that for every transition from $\underline{x}(k-1)$ to $\underline{x}(k)$, $\underline{v}(k)$ is uniquely determined. We will now show that something similar occurs with the intertransition times.

Consider two connected transitions, t_1 and t_2 , of Figure 4. There are two possible definitions of the intertransition times, given by equations (14) and (15). If equation (14) applies, then

$$\begin{aligned} v_{12}(k) &= x_2(k) - x_r(k) - [x_1(k) - x_r(k)] \\ &= w_2(k) - w_1(k) \end{aligned} \quad (21)$$

If equation (15) holds, then we can show by representing the equation in (max, +) form, using firing time state equation (4), and remembering the definition of relative firing times, that

$$v_{12}(k) = w_2(k) \odot \{w_1(k-1) \odot [a_r^T[z(k)] \odot w(k-1)]^{-1}\}^{-1} \quad (22)$$

Thus, the intertransition times can be calculated by knowing $w(k-1)$, $w(k)$, and $z(k)$.

So, analogous to the situation we had with the firing times process, for every state transition in the $\{w(k), k \in N\}$ Markov Process, the intertransition time vector $\underline{v}(k)$ is uniquely determined. This implies that the state probability distribution for the intertransition times, $F(\underline{V})$, converges since, as we saw previously, the state probability distribution

for the relative firing times converges.

IX. PERFORMANCE MEASURES

In the previous section, we showed that the intertransition times Markov process has an ergodic probability distribution. Now we will show how to evaluate the performance measures of interest.

The easiest performance measure to calculate is γ_j , the average time a token spends in place p_j . It is simply the average of the intertransition time corresponding to place p_j . Thus, $\underline{\gamma}$, the vector of average holding times is

$$\underline{\gamma} = \sum_{\underline{V}} \underline{V} F(\underline{V}) \quad (23)$$

To find λ_j , the average firing rate of transition t_j , we must think about the way live and safe DFSTPNs operate. Due to the liveness and safeness assumptions, every transition is contained in a directed circuit containing one token. Thus, the average time between firings of transition t_j , which we denote by D_j , is given by the sum of the average intertransition times around the directed circuit containing one token. By applying Little's formula [7] to the mentioned circuit, we conclude that the average firing rate is given by

$$\lambda_j = \frac{1}{D_j} \quad (24)$$

To make matters simpler, the consistency result for strongly connected DFSTPNs [14] implies that the average firing rate for all transitions is equal. Therefore, all we have to do to find the average firing rate is sum

the average intertransition times along any directed circuit and invert.

The last performance measure of interest for DFSTPNs is N_j the average number of tokens place p_j contains, for all places. This statistic can be computed by applying Little's formula to each place. Doing this, we obtain

$$N_j = \lambda \gamma_j \tag{25}$$

where λ is the average rate of arrival of tokens to p_j and γ_j is the average time a token spends in p_j .

X. EXAMPLE

Consider the DFSTPN of Figure 5, with the processing time $z_1(k)$ being 1 or 3 with equal probability, $z_4(k) = 2$ and $z_2(k) = z_3(k) = 0$. We assume these processing times are stationary and independent from place to place and over stages. The intertransition times are given by equations (18) through (21). As can be noted from these equations, it is only necessary to iterate between $p_v(k)(V)$ and $p_v(k)(V)$. Doing this, we obtain the limiting distributions shown in Figure 6. From these distributions, we find that

$$\underline{Y}^T = [13/6 \quad 7/6 \quad 1 \quad 13/6],$$

$$\lambda = \frac{6}{13}, \text{ and}$$

$$\underline{N}^T = [1 \quad \frac{7}{13} \quad \frac{6}{13} \quad 1],$$

XI. CONCLUSIONS

We have shown, in this paper, that the performance of a DFSTPN satisfying the assumptions given in the paper can be analyzed by recursively calculating the probability distributions for the intertransition times, and repeating these calculations until convergence. The solution procedure thus extends the use of state equations started by Cohen, et. al. [2], Cohen, et. al. [3], and Dubois and Stecke [6] to the stochastic case. The solution procedure also benefits from the work of the mentioned authors. Their relatively efficient algorithms can be used to find a recurrent state, and thus avoid wasting any computational effort on transient states.

Of particular note is the development of the Unfolded DFSTPN. These unfolded nets rid the problem of the complications arising from directed circuits and provide intuition for the characterization of the different processes. Moreover, it is a concept that is likely to extend to broader subclasses of STPNs.

Finding the steady state performance measures using the intertransition times state equations is more flexible than previous approaches for solving STPNs, which consisted of transforming the problem into an equivalent Markov Chain. The mentioned state equations are valid regardless of what processing time distributions are assigned. Thus, one could, for example, solve the problem by starting with gross approximations to the given processing time distributions, perform the recursive computations until convergence, then use the results to start a new set of recursive computations using better approximations to the processing times distribution⁶. Such a procedure is not possible using previous approaches

since, every time a processing time distribution is changed, a new Markov Chain has to be constructed. It is also hoped that solving these state equations recursively is numerically more stable than solving the large sets of linear equations associated with the Markov Chains.

There are several possible areas of future research. Extending the analysis to broader classes of problems is one. The efficient numerical implementation of the resulting algorithms, an issue not addressed in this paper, is another. On a longer term, one can think of performing a sensitivity analysis, doing aggregation studies, hierarchical decomposition, etc. STPNs seem to be an ideal area in which to study many topics of interest in system theory.

⁶This idea is due to Ms. E.L. Hahne, whose cooperation is greatly appreciated.

REFERENCES

- [1] E. Cinlar, Introduction to Stochastic Process, Prentice-Hall, Inc., Englewood Cliffs, N.J. (1975).
- [2] G. Cohen, D. Dubois, J.P. Quadiat, and M. Viot, "A Linear-System-Theoretic View of Discrete Event Processes and Its use for Performance Evaluation in Manufacturing", IEEE Transactions on Automatic Control, Vol. AC-30, No. 3 (March 1985), pp. 210-220.
- [3] G. Cohen, P. Moller, J.P. Quadrat, and M. Viot, "Linear System Theory for Discrete Event Systems", Proceedings of 23rd Conference on Decision and Control, Vol. 1 (December 1984), pp. 539-544.
- [4] F. Commoner, A.W. Holt, S. Even, and A. Pnueli, "Marked Directed Graphs", Journal of Computer and System Sciences, Vol. 5 (1971), pp. 511-523.
- [5] M. Diaz, "Modeling and Analysis of Communication and Cooperation Protocols using Petri Net Based Model", Computer Networks, Vol. 6, No. 6 (December 1982), pp. 419-441.
- [6] D. Dubois and K.E. Stecke, "Using Petri Nets to Represent Production Processes", Proc. of the 22nd IEEE C.D.C., Vol. 3 (December 1983), pp. 1062-1067.
- [7] L. Kleinrock, Queueing Systems - Vol. 1: Theory, John Wiley and Sons, New York, N.Y. (1975)
- [8] M.A. Marsan, G. Conte, and G. Balbo, "A Class of Generalized Stochastic Petri Nets for the Performance Evaluation of Multiprocessor Systems", ACM Transactions on Computer Systems, Vol. 2, No. 2 (May 1984), pp. 93-122.
- [9] M.K. Molloy, "Performance Analysis Using Stochastic Petri Nets", IEEE Trans. on Computer, Vol. C-31, No. 9 (September 1982), pp. 913-917.
- [10] T. Murata, "Circuit Theoretic Analysis and Synthesis of Marked Graphs", IEEE Trans. on Circuits and Systems, Vol. CAS-24, No. 7, (July 1977), pp. 400-405.
- [11] J.L. Peterson, Petri Net Theory and the Modeling of Systems, Prentice-Hall, Englewood Cliffs, N.J. (1981)
- [12] C.V. Ramamoorthy and G.S. Ho, "Performance Evaluation of Asynchronous Concurrent Systems Using Petri Nets," IEEE Trans on Software Eng., Vol. SE-6, No. 5 (September 1980), pp. 440-449.
- [13] C. Ramchandani, "Analysis of Asynchronous Concurrent Systems by

Timed Petri Nets", Technical Report 120, Lab. for Computer Science, M.I.T., Cambridge, MA. (1974).

- [14] J. Sifakis, "Use of Petri Nets and Preliminary Performance Evaluation", The 7th Annual Symposium on Computer Architecture - Conference Proceedings, (May 6-8, 1980), pp. 88-96.
- [15] W.M. Zuberek, "Timed Petri Nets and Preliminary Performance Evaluation", The 7th Annual Symposium on Computer Architecture - Conference Proceedings (May 6-8, 1980), pp. 88-96.




Figure 1. Stochastic Timed Petri Net Example
(without processing times)




Figure 2. Stochastic Timed Decision-Free Petri Net Example
(without processing times)




Figure 3. Unfolded STDFPN that corresponds to STDFPN of Figure 2.





Figure 4. STDFPN for Firing Time State Equation derivation



$$x_1(k) = \max \{ z_1(k) + x_1(k-1), z_3(k) + x_2(k-1) \}$$

$$x_2(k) = \max \{ z_2(k) + x_1(k), z_4(k) + x_2(k-1) \}$$

Figure 5. STDFPN Example with Firing Time State Equations

$k \rightarrow \infty$





Figure 6. Limiting Probability Distributions for Numerical Example.