A GENERALIZED TIMED PETRI NET MODEL
FOR PERFORMANCE ANALYSIS

by

Mark A. Holliday and Mary K. Vernon

Computer Sciences Technical Report #593

May 1985

A Generalized Timed Petri Net Model for Performance Analysis

Mark A. Holliday and Mary K. Vernon

Computer Sciences Department
University of Wisconsin — Madison
Madison, WI 53706

Abstract

We have developed a Generalized Timed Petri Net (GTPN) model for evaluating the performance of computer systems. Our model is a generalization of the TPN model proposed by Zuberek [1] and extended by Razouk and Phelps [2]. In this paper, we define the GTPN model and present how performance estimates are obtained from the GTPN. We demonstrate the use of our automated GTPN analysis techniques on the Dining Philosophers example. This example violates restrictions made in the earlier TPN models. Finally, we compare the GTPN with the Stochastic Petri Net (SPN) models. We show that the GTPN model has more modeling power than the existing SPN models. The GTPN provides an efficient, easily used method of obtaining accurate performance estimates for models of computer systems which include both deterministic and geometric holding times.

1. Introduction

We present a Generalized Timed Petri Net (GTPN) model for evaluating the performance of computer systems. The GTPN model is an efficient, easily used tool for calculating exact performance estimates for many models of computer systems. For example, in [3] we derive exact performance estimates of multiprocessor memory and bus interference for fixed memory cycle times. Only approximate solutions existed previously for many of the models we studied.

Petri Nets are a graph model of computation [4]. Modifying Petri Nets so that time is represented has recently been an active research area. The goal of these models is to analyze system performance as an extension of the reachability analysis. Our model is a generalization of the Timed Petri Net (TPN) model proposed by Zuberek [1] and extended by Razouk and Phelps [2]. The TPN model associates firing frequencies and deterministic firing times with each transition in the net. Both Zuberek and Razouk and Phelps restrict the allowed nets in order to propose algorithms for building the net's reachability graph. They further restrict the allowed nets in order to analyze the net's reachability graph. Our generalization, the GTPN model, removes all of these restrictions. We present efficient algorithms for building the reachability graph of an arbitrary bounded net and we show how to analyze that reachability graph in the general case.

For the purpose of performance analysis we view the GTPN as a stochastic process. The time-in-state is a deterministic function for each state of the net. However, a probability distribution is defined over the possible next states based on the firing frequencies. The GTPN analyzer automatically generates the associated discrete parameter, embedded Markov Chain and calculates performance estimates.

Stochastic Petri Nets (SPNs) are an alternate method of representing time in Petri Nets for the purpose of performance analysis. Continuous-time SPN models were proposed by Natkin[5], Symons[6], and Molloy [7,8,9], and generalized by Marsan, Balbo, and Conte [10] and Dugan, Trivedi, Geist, and Nicola [11]. The generalized continuous-time SPN (GSPN) model associates an instantaneous or exponentially distributed firing time with each transition. Molloy [7,9] has also

proposed a discrete-time SPN model with transition firing times that are geometrically distributed.

The discrete-time SPN model is especially interesting because it can represent deterministic firing times, thus bridging the gap between the SPN and TPN models. However, the representation of deterministic firing times has two restrictions in the discrete-time SPN: 1) a firing time must be a non-zero multiple of some unit time step, and 2) all conflicting actions having deterministic delays must be equally likely.

The GTPN can represent geometric holding times, and thus also bridges the gap between the TPN and SPN models. Furthermore, deterministic firing times can be any nonnegative real value, including zero, and we can assign arbitrary next-state probabilities to conflicting transitions. The GTPN is thus more powerful than existing SPN models. We examine this issue in more detail in this paper.

This paper is an extended version of [12]. In section 2, we describe untimed Petri Nets and the previous TPN models. In Section 3 we describe the GTPN model and in Section 4 we show how the GTPN model is analyzed. In section 5, we demonstrate the use of our automated GTPN analysis techniques on the Dining Philosophers model [13], which violates restrictions in the earlier TPN models. In Section 6 we compare the GTPN with the discrete-time and continuous-time SPN models. Section 7 contains the conclusions of this work and suggestions for future research.

2. Untimed Petri Nets and Previous TPN Models

In this section we describe untimed Petri Nets and previous TPN models which provide a foundation for the GTPN in Section 3. Section 2.1. describes untimed Petri Nets. A more thorough introduction to untimed Petri Nets can be found in Peterson[13]. Section 2.2 reviews the work of Zuberek and Razouk and Phelps.

2.1. Untimed Petri Nets

Untimed Petri Nets (PNs) contain *places P*, *transitions T*, and *arcs A*. The arcs are directed and can only connect transitions to places and places to transitions. If an arc exists from a place to a transition, then the place is an *input place* for that transition. If an arc exists from a transition

to a place, then the place is an *output place* for that transition. Places may contain *tokens*. The *state* of a PN is defined by the number of tokens in each place and is represented by a vector M called the *marking* vector. $M[i]$ is the number of tokens in the ith place. A transition is *enabled* if each of its input places contains at least as many tokens as there are arcs from the place to the transition. The tokens in the input places which enable a transition are its *enabling* tokens. An enabled transition can *fire*. A transition *fires* by: 1) removing all of its enabling tokens from its input places, and 2) placing on each of its output places one token for each arc from the transition to that output place. Each firing of a transition changes the assignment of tokens to places and thus creates a new state. The *reachability set* of a PN and a given initial state is the set of all states that can be reached from that initial state via a sequence of transition firings. The *reachability graph* associated with a reachability set can be constructed as follows. Represent each state by a vertex and place a directed edge from vertex v_1 to vertex v_2 if the state v_2 can result from firing some transition enabled in state v_1.

Petri Nets are often illustrated graphically. Circles represent the places. Black dots in the circles represent the tokens. Bars represent the transitions. Figure 3.1 can be considered an example of an untimed Petri net by simply ignoring the vectors next to each transition.

2.2. Previous TPN Models

Ramachandani [14] was the first to introduce a fixed firing time with each transition in a Petri net. Ramamoorthy and Ho [15], Zuberek [1], and Razouk and Phelps [2] are three more recent studies that use a single fixed firing time. Our work is based on the TPN model of Zuberek and Razouk and Phelps.

The TPN model [1][2] is a Petri net which has been augmented to include a set of firing durations (D), a set of firing frequencies (F), and a set of named resources (R). Each set is associated with the transitions in the net. The model is thus formally defined as follows:

$$TPN = (P, T, A, M', D, F, R)$$

where

3

$P = \{p_1, p_2, \ldots, p_n\}$ (places)

$T = \{t_1, t_2, \ldots, t_m\}$ (transitions)

$A \subseteq \{P \times T\} \cup \{T \times P\}$ (directed arcs)

$M' = \{m_1, m_2, \ldots, m_n\}$ (initial marking)

$D = \{d_1, d_2, \ldots, d_m\}$ (firing durations)

$F = \{f_1, f_2, \ldots, f_m\}$ (firing frequencies)

$R = \{r_1, r_2, \ldots, r_m\}$ (resources)

The state of a TPN is defined differently than in untimed Petri Nets because firing a transition is not an atomic operation. A transition has an associated deterministic firing duration. There is a *start firing*, and an *end firing* event. In between the firing is *in progress*. The removal of tokens from a transition's input places occurs at start firing. The placement of tokens on a transition's output places occurs at end firing. While the firing of a transition is in progress, the time to end firing, called the *remaining firing time* (RFT), decreases from the firing duration to zero (without causing a change in the marking of the net). Because firings can be in progress when a marking change occurs, a state is only partially defined by the distribution of tokens. A state must also include the RFT of each firing in progress. A state is thus a marking vector and a set of RFT's.

Also unlike an untimed Petri net, the next state is not generated by a single start firing or end firing event. Instead it is generated by a *set* of start firings or a *set* of end firings which occur simultaneously. Given a particular state, the basic rule for finding the possible next states is straightforward. Find how many *enablings* of each transition exist. This is determined by the number of tokens on its input places. Find the *maximal* sets [1] that can start firing simultaneously. Each maximal defines a next state. The time spent in the original state is zero. The RFT vectors of the transitions which just started firing are set to their transition's duration. If there are no enablings, but there are some firings in progress, then the next state is generated by the end firing of all transitions with the smallest RFT (Tmin). The time-in-state value in this case is Tmin. If

---

[1] A set with property α is a *maximal* set with property α if it is not a proper subset of any other set with property α.

there are no enablings and no firings in progress, then the net remains in the current state forever.

The rule that next states are generated by sets of events that occur simultaneously, is not strictly necessary. The advantage of having it is that the state spaces generated are dramatically smaller. The disadvantage is that the algorithms for building the state space are more complicated.

The frequencies are used in assigning probabilities to next states when building the net's reachability graph.

Zuberek suggested that the reachability graph of the Timed Petri Net be viewed as a Markov Chain and that performance measures be computed using standard techniques for analyzing the Markov Chain's long run behavior. Extensions of his work, however, are desirable in two areas. One, he only proposed a method for constructing a net's reachability graph for a restricted class of nets: nets that are *safe* and *free choice*. A net with a given initial state is said to be *safe* if, for every state in the reachability set, no place has more than one token. A place is a *decision place* if it is an input place to more than one transition. A net is *free choice* if every transition that has an input place that is a decision place has no other input places. Two, even for safe and free choice nets, the structure of the reachability graph (i.e. the Markov Chain) may be such that Zuberek's approach gives incorrect values for the performance measures. The states in a discrete time Markov Chain can be divided into classes. A set of classes, called *recurrent* classes, is important because in the long run the model will reach and stay in one of these classes. Zuberek's approach gives correct values only when there is exactly one recurrent class.

Razouk and Phelps [2] extend Zuberek's work in the first of the two areas above. They allow a superset of the class of safe, free choice nets. Two or more transitions are said to be in the same *conflict set* if their sets of input places intersect. Two conflict sets *overlap* if at least one transition is in both. Razouk and Phelps make the restrictions that conflict sets do not overlap and that all transitions in a conflict set are *mutually disabling*, i.e. firing of one, disables all the others. They maintain Zuberek's restriction in the second area (they call this, requiring a *cyclic* net).

Razouk and Phelps also introduce the concept *resources*, originally proposed in E-Nets[16]. A

5

resource can be associated with one or more transitions. Whenever one of those transitions is firing, the resource is in *use*. If more than one of these transitions is firing simultaneously, the resource has several *usages* occurring. By building and analyzing the net's reachability graph we can find the average number of uses of a resource over time. This average, if properly implemented and interpreted, can be used to obtain a variety of meaningful performance estimates.

3. The GTPN Model

The GTPN model extends the models of Zuberek and Razouk and Phelps by: 1) removing all restrictions on the net except the obvious one that the state space be finite, and 2) computing correct performance estimates for any reachability graph (i.e. an arbitrary embedded discrete parameter, finite state Markov Chain). We also allow the firing duration to be an arbitrary real number (the non-integer case is not discussed by Zuberek or Razouk and Phelps).

A third extension we have found useful involves firing durations and frequencies. In the models of Zuberek and Razouk and Phelps, the duration and frequency are state-independent constants. Consequently, in the GTPN model a transition's firing duration and frequency can be expressions containing immediate values (real and integer), names of places (a place name stands for the number of tokens in that place in the current state), and arithmetic, relational, and logical operators. These state dependent durations and frequencies become deterministic values when used to determine time-in-state and next state probabilities for each state in the reachability graph.

Besides a firing duration, frequency, and set of resources, a GTPN transition has a flag associated with it that is used in computing the next state probabilities as described in section 3.2.

Figure 3.1. shows an example GTPN net, including the initial state distribution of tokens. Each place and transition is labelled. Each transition has, from left to right, its firing duration expression, its frequency expression, its flag, and its list of resources. This example models users at terminals, who with a geometric think time generate requests for a server. There is one token on place P1 for each user. Transitions T1 and T2 implement the think time. Transition T3 implements a load-dependent server with a firing duration that depends on the number of tokens on P2.
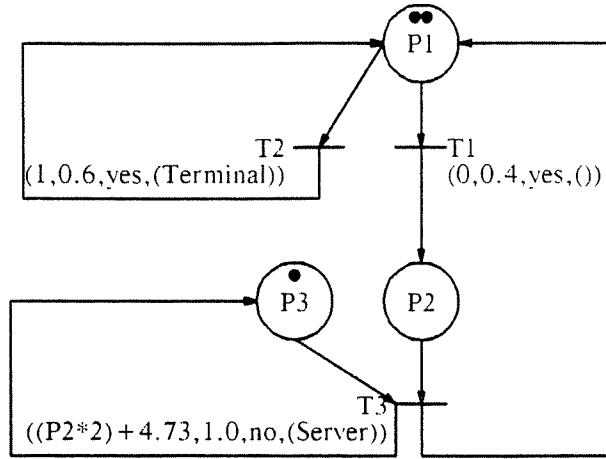
Figure 3.1. Example of a GTPN net

Figure 3.2. Reachability Graph for example

Table 3.1. Reachable States for example

States	P1	P2	P3	RFT Set	Resources
0	1	0	1	{ }	{ }
1	0	0	1	{(T1,0.0)}	{ }
2	0	0	1	{(T2,1.0)}	{Terminal(1)}
3	0	1	1	{ }	{ }
4	0	0	0	{(T3,6.73)}	{Server(1)}

In Figure 3.2 and Table 3.1 we show the reachability graph for the simple GTPN in Figure 3.1 assuming only one user. The labels on the edges of the graph are the next state probabilities. The labels on the vertices of the graph are the values for time-in-state. The marking vectors are shown in the table. The RFT sets are shown as a list of pairs with one pair per in progress firing of a

7

transition. The first component of each pair is the name of the transition. The second component is the remaining firing time. The resources used and their number of uses are also shown in the table.

The Razouk and Phelps' TPN model does not allow multiple tokens on place P1. Allowing such nets complicates constructing the reachability graph. An overview of our reachability algorithm, which handles these complications, is in Figure 3.3. The TimeInState and ResUsages (number of usages of a resource) functions are used in the performance analysis as described in Section 4. The algorithm has two complex parts: 1) finding the next states when the next states are due to maximal sets of transition enablings which start firing together, and 2) assigning probabilities to next states. These two parts are discussed in the Sections 3.1 and 3.2.

3.1. Finding Maximals

Our first point is somewhat discouraging. In an arbitrary state in a net the number of maximals in the worst case is exponential in the number of enablings.

Theorem 1. *Consider a state S in a GTPN with n enablings. The number of maximals is* $\Omega(2 \frac{2^n}{\sqrt{2\pi n}})$.

pf. First we will construct a state in a net such that the number of maximals is $\binom{n}{n/2}$. Consider the case where n is even (The n is odd case is similar.). Consider the net with n transitions and one place P where the one place is an input place for all the transitions. Assume each transition needs one token in the input place to have one enabling. Let the state have $n/2$ tokens on the input place P. Then there are $\binom{n}{n/2}$ maximals. Thus, the number of maximals is $\Omega\binom{n}{n/2}$. The observation that by Stirling's approximation, $n! = \sqrt{2\pi n}(n/e)^n(1 + O(1/n))$, completes the proof. ∎

This result should not be given too much weight. In practice, we find that the number of maximals is far less than exponential in the number of enablings. Theorem 1 does, however, point out that the space of potential maximals is large. Consequently, an algorithm for finding the maximals must be carefully thought out in order to prevent poor performance when the actual number of maximals is small. The algorithm described below meets this criterion. When we profiled

```
X.State ← Initial State; X.Class ← Frontier
while at least one Frontier state, Y do begin
      if Y is a duplicate of an Interior state Z then
            Y.Class ← Duplicate
      else begin
            Find the set of enablings in Y
            If no enablings and the RFT set is empty then
                  Y.Class ← Terminal
            else if any enablings then begin
                  Find the set of maximals of enablings
                  Compute the probability of each maximal
                  For each maximal M create a new state Z from Y
                        Remove tokens from the input places of
                              transitions that have enablings in M
                        Add a firing, $f_t$, to the RFT set for each
                              enabling in M
                        Set the RFT of each added firing, $f_t$, to the
                              firing duration of transition t
                        Z becomes a child of Y; Z. Class ← Frontier
                  for all resources ResUsages Y ← count uses
                  TimeInState[Y] ← 0; Y.Class ← Interior end
            else begin
                  Let Tmin be the smallest RFT in Y
                  Create state Z from Y by subtracting Tmin from
                        each RFT in Y
                  For each firing $f_t$ whose RFT = 0 in Z do
                        Add tokens to the output places of transition t
                        Remove $f_t$ from the RFT set
                  Z becomes a child of Y; TimeInState[Y] ← Tmin
                  for all resources ResUsages[Y] ← count uses
                  Z.Class ← Frontier; Y.Class ← Interior end
      end
```

Figure 3.3. Overall State Space Algorithm

our GTPN tool, the percentage of total program time taken by this algorithm was less than 5% for

the analysis of large nets [2].

Our algorithm consists of two independent subalgorithms *Partition* and *FindMax*. The Partition algorithm partitions the set of enablings into *Generalized Conflict Sets*, such that maximals of the partitioned sets can be efficiently combined to generate all the maximals for the original set of

---

[2] we profiled our program using *gprof* under 4.2 bsd UNIX[TM] running on a VAX-11/780[TM]

9

enablings. The Partition algorithm does not specify how the maximals for each partition member are found. Maximals are found for each member of the partition by FindMax.

3.1.1. The Partition Algorithm

The Partition algorithm has two parts. The first part constructs a static partition of the set of transitions, T. Define the *directly-conflicts-with* relation on T as follows. For all t_1 and t_2 in T, t_1 directly-conflicts-with t_2 if the set of input places for t_1 intersects the set of input places for t_2. Define the *conflicts-with* relation on T as the transitive closure of the directly-conflicts-with relation. The conflicts-with relation is clearly reflexive, symmetric, and transitive, so it is an equivalence relation on T. The partition of T induced by conflicts-with is denoted by $\{GCS[i] | 1 \leq i \leq N\}$, where $GCS[i]$ is the ith member of the partition and N is the size of the partition (GCS stands for *generalized conflict set*). Note that a transition which does not share any input places with any other transitions forms a GCS of size one.

Part two of the Partition algorithm uses the conflicts-with partitioning of T to partition the set of enablings, $Enablings(S)$, for each state S. The desired partition of $Enablings(S)$ is

$$\{EGCS[S, i] = Enablings(S) \cap GCS[i] | \text{partition i}\}.$$

At this point, FindMax is applied to each partition member to find its *local maximals*. Let $Maximals[Enablings(S)]$ be the set of maximals over all the enablings. We have constructed our partition such that $Maximals[Enablings(S)]$ is simply the Cartesian product of the local maximals.

$$Maximals[Enablings(S)] = FindMax(S, EGCS[1]) \times \quad \times FindMax(S, EGCS[N])$$

3.1.2 The FindMax Algorithm

We want to find the local maximals for a generalized conflict set, G, in a given state S. Any subset of $EGCS[S, G]$ is a potential local maximal (for brevity's sake we will call a local maximal, a maximal, in this section) The power set, $P(EGCS[S, G])$, unfortunately, can be a large search space. Our goal is to minimize the number of members of this power set that we examine. Note that set inclusion defines a partial order on $P(EGCS[S, G])$ which, in turn, induces a directed acyclic graph

10

This graph has one root which represents the set $EGCS;S.G$ itself. The FindMax algorithm does a breadth-first search of this graph searching for vertices that are maximals. The traversal is implemented in the standard way using a queue. Initially, the root vertex is the only entry on the queue.

Searching the graph breadth-first causes the order in which vertices are examined to have an important property: if 1) the set of enablings, E_1, represented by a vertex can fire together, and 2) E_1 is not a subset of any maximal already found, then E_1 is a maximal. In other words, it is impossible that some vertex examined later will have a set of enablings, E_2, that can all fire together and for which E_1 is a subset. Thus, to find the maximals we just do the breadth-first search, checking each vertex to see if it satisfies properties 1) and 2).

Three methods are used to avoid searching the entire graph. One, if a vertex's set of enablings, E_1, can all fire together, then none of its descendents needs to be examined (so its children are not added to the queue). To see this, note that either E_1 is a maximal or it is a subset of a maximal that has already been found. Two, a pointer is used to ensure that vertices are never examined more than once. Note that the graph is not a tree. Thus, a naive breadth-first search would cause vertices to be examined multiple times. Three, the pointer used in method two is used to implement a heuristic for pruning subtrees.

Methods two and three are based on a pointer, V_p, associated with each vertex, V, in the graph. Consider the set of enablings, E_1, represented by V to be described by a vector of nonnegative integers. The ith component of the vector gives the number of enablings of the ith transition. V_p points at one of the components in this vector. The set of transitions to the left of V_p is V_L. The set of transitions to the right of V_p and including the transition pointed to by V_p is V_R.

The pointer V_p can be used to ensure that a vertex is only examined once. The method is as follows. Set the V_p of the root vertex to point to the leftmost transition. For each parent vertex, subtracting one enabling from the transitions in V_R defines a child vertex. For each of these child vertices, set its V_p to point to the transition that was decremented. It can be shown inductively that

at each level in the graph, the sets V_L are distinct for each vertex. Consequently, any child resulting from decrementing V_R of one parent can never be the same as a child resulting from decrementing V_R of another parent. Note, also, that the use of V_l does not cause any state to not be examined that should be examined. For each state, c, that should be examined, there is a vertex, p, whose set of enablings over the transitions in V_L matches c's set of enablings over the same transitions. Consequently, c will become a child of p.

Method three is a simple optimization made possible by the existence of V_l. If the enablings in V_L cannot fire together, then it is irrelevant what enablings are subtracted from V_R. Consequently, none of the child vertices resulting from decrementing V_R need be added to the queue.

Partition and FindMax are used to calculate all maximal sets of enabled transitions which can start firing together. Next state transition probabilities must be assigned to these maximal sets. This is discussed in the next section.

3.2. Computing Probabilities

To interpret the reachability graph as a Markov Chain we need to assign probabilities to next states. In the nontrivial case, we need to assign a probability to each maximal set of transitions that can start firing together. From the previous section we know that a maximal is the union of a set of *independent* local maximals, one from each GCS. Thus a reasonable probability for the maximal is the product of the probabilities for the local maximals. Suppose that $LocalMax[k, i]$ is the kth local maximal of the ith generalized conflict set. Suppose that there are N generalized conflict sets and the jth global maximal is the union over all GCS's of the j_ith local maximal of the ith generalized conflict set. Then

$$Pr\{Maximal[j]\} = \prod_{|i:i=1 \quad N|} Pr\{LocalMax[j_i, i]\}$$

In order to compute the probability of a local maximal, we take the product of the frequencies of all the enablings in the local maximal. This is multiplied by a number $NumComb$ discussed below. Then for each local maximal, we normalize this product by dividing it by the sum of the products over all local maximals. More formally, suppose the ith GCS has M local maximals

and the frequency expression for the kth enabling in the ith GCS is f_k. Then our formula for $Pr\{LocalMax[j_i, i]\}$ is

$$Pr\{LocalMax[j_i, i]\} = \frac{NumComb[LocalMax[j_i, i]] * \prod\limits_{\{k:k \in LocalMax[j_i, i]\}} f_k}{\sum\limits_{\{m:m=1...M\}} NumComb[LocalMax[m, i]] * \prod\limits_{\{k:k \in LocalMax[m, i]\}} f_k}$$

$NumComb$ means *number of combinations* and is a combinatoric value associated with each local maximal. This value is defined as the number of ways tokens can be removed from input places in order to implement that local maximal. As a simple example, suppose the local maximal consists of one enabling of one transition with one input place, two arcs connect the place to the transition, and the input place has three tokens. In this case, the combinatoric value $NumComb$ is $\binom{3}{2}$.

Computing $NumComb[LocalMax]$ is done by decomposing it first on the transitions in the local maximal and second on the input places for the given transition. The number of ways that the tokens can be removed from the given input place by the given transition is a binomial coefficient. As we consider each transition, the value in the upper position of this binomial coefficient is changed to reflect any removals made by transitions considered earlier.

More formally, consider the local maximal, $LMax$, of the ith generalized conflict set. Let $InputPlace[t]$ be the set of input places for transition t. Let $Enab[t]$ be the number of enablings of transition t in this maximal. Let $TokNeeded[t, p]$ be the number of tokens needed from input place p by one enabling of transition t. Note that $Enab[t] * TokNeeded[t, p]$ is the number of tokens needed from place p by transition t in this maximal. Let $TokLeft[p]$ initially be the number of tokens on input place p in the parent state. After looking at an input place $TokLeft[p]$ is updated to reflect the start firing of all the enablings of transition t. With this notation we have:

$$NumComb[LMax] = \prod_{\{t:t \in LMax\}} \prod_{\{p:p \in InputPlace[t]\}} \binom{TokLeft[t, p]}{Enab[t] * TokNeeded[t, p]}$$

From practical experience it appears that in some cases it is reasonable to use this combinatoric value when assigning probabilities to local maximals. A boolean flag associated with each transition specifies whether this should be done. Only if the flag is *yes* for all transitions in the maximal, is $NumComb$ used. Formula 1 above is for the case where $NumComb$ is used. The formula in the alternative case is the same except that $NumComb = 1$ for all terms.

13

The motivation for our method of assigning probabilities to local maximals is that it assigns the right probabilities in the important case where all the enablings in a local maximal are independent events. In the case where there are dependencies between the enablings it is difficult to envision a single formula that will always generate the "right" probabilities. This case motivated our introduction of state-dependent frequency expressions. Such frequency expressions can specify what probabilities maximals should have in different markings. Allowing state-dependent frequency expressions also allows the possibility that in some states a transition's frequency expression may evaluate to zero even though it has one or more enablings. We remove these enablings from the set of considered enablings before finding the local maximals.

4. GTPN Performance Analysis

For the purpose of performance analysis, we view the GTPN as a stochastic process. The time-in-state is a deterministic function, TimeInState, of the state. Nevertheless, the process is stochastic because of the probability distribution over the possible next states. Since the time-in-state can be an arbitrary real number, the process is a continuous time stochastic process. The parameter set is described in Section 4.5. The states of the stochastic process are divided into classes. In the long run, with probability one, the process will reach and stay in one of the set of classes called *recurrent* classes [3]. Consequently, the *long run fraction of time spent in each state* depends on which recurrent class the process reaches in the long run. For each recurrent class the long run fraction of time spent in each state forms a probability distribution over the states. Thus, there is a vector of long run probability distributions with one component for each recurrent class. In addition, we can compute the probability (the *absorption* probability) of reaching each recurrent class in the long run. These absorption probabilities allow us to assign relative weights to the components of the vector of long run probability distributions.

The number of usages of a given resource is also deterministic for a given state; it is a function ResUsages of the state. Consequently, ResUsages, being a function of a random variable, is a

---

[3] Note that a terminal state in the reachability graph is a recurrent class due to the self-loop we added (see Figure 3.3).

random variable. A performance estimate for a resource is a vector with one component for each recurrent class. The value for recurrent class R's vector component is the long run expectation of that resource's ResUsages random variable with respect to R's long run probability distribution. In other words, it is the weighted sum of the long run fractions of time spent in each state given that in the long run the process is in class R The weight of a state is the number of resource usages in that state.

Our approach to computing performance estimates uses the key observation [4] that the times at which state changes occur form an embedded, discrete time, finite state Markov Chain. Consequently, our approach has four parts: 1) building the Markov Chain (i.e. reachability graph; described in Section 3), 2) aggregating the states in order to reduce the size of the state space (Our aggregation rule is: Any state S_2 can be aggregated with its parent state, S_1, if and only if S_1 is S_2's only parent and S_2 is S_1's only child.), 3) analyzing the Markov Chain, and 4) computing resource usage expectations in the original stochastic process.

Part 3, the Markov Chain analysis, has three steps: a) finding the chain's recurrent classes, b) finding the absorption probability for each recurrent class, and c) finding, for each recurrent class, the long run fraction of visits to each state. These steps are discussed, respectively, in Sections 4.1, 4.2, and 4.3.

Part 4 has two steps: a) for each recurrent class R, computing the long run fraction of time spent in each state (the TimeInState function and the long run fraction of visits to each state are used to do this), b) for each recurrent class R, use the ResUsages functions and the long run fraction of time spent in each state to find the long run expected number of usages of each resource These two steps are discussed in Section 4.4

As mentioned above, these long run expectation are the performance estimate (given that the process is in class R in the long run). For each resource we thus have a vector of performance estimates. If desired, the vector components could be weighted by the absorption probabilities to

---

[4] The GSPN model uses a similar observation.

give a single performance estimate.

In Section 4.5 we give a more precise definition of the parameter set of the GTPN stochastic process.

4.1. Finding Recurrent Classes

In order to find the recurrent classes we need to first define a *recurrent* class. Recall that $\forall n$ $P_{ij} = Pr\{X_{n+1} = j | X_n = i\}$. $P = |P_{ij}|$ is the *one-step transition probability matrix*. $P^{(n)}$, the *n-step transition probability matrix* is defined similarly. $f_{ii}^{(n)}$ is the probability that, starting from state i, the first return to state i occurs at the nth transition. A state is *recurrent* if $\sum_{n=0}^{\infty} f_{ii}^{(n)} = 1$. In other words, a state is recurrent if and only if, after the process starts from state i, with probability one the process returns to state i in a finite length of time. A state is *transient* if it is not recurrent. State j is said to be *accessible* from state i if $P_{ij}^{(n)} > 0$ for some integer $n \geq 0$. Two states i and j that are each accessible to the other, are said to *communicate*.

Note the following facts. *Accessibility* defines a partial order on the states. This partial order implies a directed graph with the vertices being the states. *Communication* is an equivalence relation on the states. Thus it partitions the states into subsets called *classes*. The *communication* classes are the strongly connected components of the accessibility graph. These strongly connected components form a directed acyclic graph, DAG.

All the states in a class are recurrent or none are, so we can speak of *recurrent* classes and *transient* classes. In the case of a finite state space [5], the *recurrent* classes are the leaves of the DAG[17] of strongly connected components. The interior nodes of the DAG are the *transient* classes.

Given that the recurrent classes are the leaves of this DAG, the algorithm to find the recurrent classes is immediate. The reachability graph of the GTPN is the accessibility graph. Create the DAG by finding the strongly connected components of the accessibility graph. We do this using Tarjan's $O(nlogn)$ algorithm[18]. Then find the leaves of the DAG by a depth first search starting

---

[5] This is not true if the the state space is countably infinite. A simple counterexample is a one-dimensional, asymmetric random walk.

at the initial state.

4.2. Absorption Probabilities

In a finite state Markov Chain, if we start in a state in a transient class we will eventually reach and stay forever in one of the recurrent classes. We are said to be *absorbed* by a particular recurrent class. Computing the probability of absorption in a particular recurrent class R given a particular initial state, i, can be done using a standard technique called *first-step analysis*[19]. On the first state change, the process will move from state i to a state j that is in a transient class or in a recurrent class. If j is in class R, the future probability of being absorbed by class R is one. If j is in another recurrent class, the future probability of being absorbed by class R is zero. If j is in a transient class, then, by the memoryless property, the probability of being absorbed by class R is the same as if j were the initial state.

More formally, suppose that the states in all the transient classes are numbered $0, ..., r-1$ and consider a fixed recurrent class R and fixed initial state i. Let $U_i = Pr\{$ Absorption in class $R|X_0 = i\}$ for $0 \leq i < r$.

$$U_i = \sum_{\{j \in R\}} P_{ij} + \sum_{j=0}^{r-1} P_{ij} U_{jR}, \qquad i = 0, 1, ..., r-1$$

This equation cannot be be solved in isolation. However, if we consider all possible initial states, then we have a system of linear equations that can be solved for the U_i's.

4.3. Long Run Expected Fraction of Visits

If R is a recurrent class in a finite state space, then $\forall j \in R$ a number π_j exists such that $\forall i \in R$

$$\lim_{m \to \infty} E\left[\frac{1}{m} \sum_{k=0}^{m-1} 1_{\{X_k = j\}} | X_0 = i\right] = \lim_{m \to \infty} \frac{1}{m} \sum_{k=0}^{m-1} P_{ij}^{(k)} = \pi_j$$

The leftmost expression above is *the long run expected fraction of visits spent in state j*. $1_{\{X_k = j\}}$ is the indicator random variable that equals one when the outcome chosen is in the event $\{X_k = j\}$ and 0 otherwise.

We want to find these π_j's for each class R. We do this by noting that each recurrent class R in a finite state space has one and only one *stationary probability distribution* and the vector π_R of

its π_j's is this stationary probability distribution. This stationary probability distribution is easy to find since it is the unique solution to the set of equations

$$\pi_R = \pi_R P_R \qquad \text{and} \qquad \sum_{j \in R} \pi_j = 1$$

The matrix P_R is $P_R = \{P_{ij} | i \in R\}$. We solve this system of equations using an iterative matrix algorithm[20].

It is also true that an arbitrary Markov Chain with a finite state space has at least one stationary probability distribution over the entire state space. However, if the Markov Chain has more than one recurrent class, then any linear combination that sums to one of the stationary probability distributions of the individual recurrent classes is a stationary probability distribution of the chain as a whole.

Our approach is correct regardless of whether the recurrent classes are periodic or aperiodic. Recall that the *period* of a state is the greatest common divisor of all integers $n \geq 1$ for which $P_{ii}^{(n)} > 0$. A state is *aperiodic* if its period is 1, else is *periodic*. All the states in a class have the same period so we can refer to a class as periodic or aperiodic. Only for the states i in an aperiodic recurrent class does the limit $\lim_{n \to \infty} P_{ii}^{(n)}$ exist. However, the long run expected fractions of visits and the stationary probability distribution exist in both cases. Again we assume a finite state space.

4.4 Resource Usage Estimates

We find, for each recurrent class R, the long run expected fraction of time spent in each state. Then, we find the long run expectation of each resource with respect to each recurrent class.

Let S be the set of states. Let $RelTime(S_1)$ be the long run expected fraction of time spent in state S_1, given that the process is absorbed in class R. From the Markov Chain we know the long run expected fraction of visits to each state k, π_k, given absorption in class R. $RelTime(S_1)$ can be computed, using the TimeInState function, as follows:

$$RelTime(S_1) = \frac{\lim_{n \to \infty} \frac{1}{n} \sum_{t=0}^{n-1} E[1_{\{X(t)=S_1\}} TimeInState(X(t))]}{\lim_{n \to \infty} \frac{1}{n} \sum_{t=0}^{n-1} E[TimeInState(X(t))]}$$

$$= \frac{TimeInState(S_1)\pi_{S_1}}{\sum_{k \in S} TimeInState(k)\pi_k}$$

Recall that 1_A is the indicator random variable for the event A. To show why the last equality holds we derive its denominator. A similar derivation holds for the numerator. Let S be the set of states.

$$\lim_{n \to \infty} \frac{1}{n} \sum_{t=0}^{n-1} E[TimeInState(X(t))]$$

$$= \lim_{n \to \infty} \frac{1}{n} \sum_{t=0}^{n-1} [\sum_{k \in S} TimeInState(k)Pr\{X(t) = k\}]$$

$$= \sum_{k \in S} TimeInState(k) \lim_{n \to \infty} \frac{1}{n} \sum_{t=0}^{n-1} Pr\{X(t) = k\}$$

$$= \sum_{k \in S} TimeInState(k)\pi_k$$

To find the long run expected number of usages of each resource for each absorbing recurrent class, we simply take the expectation of the random variable ResUsages.

$$E[ResUsages] = \sum_{k \in S} ResUsages(k)Pr\{state k\} = \sum_{k \in S} ResUsages(k)RelTime(k)$$

4.5 Parameter Set of the Stochastic Process

Since firing durations can be zero in the GTPN model, a GTPN can be in two or more states at the same "time". This slightly complicates viewing a GTPN as a stochastic process. In this subsection we discuss that complication and how it can be resolved. The complication is that the right halfline cannot be the parameter set of the GTPN stochastic process. To see this, recall that a stochastic process is a family of random variables indexed by the parameter set and a random variable is a function from the sample space into the reals. If the right halfline were the parameter set, then on some sample path at some parameter t, the random variable $X(t)$ could simultaneously hold more than one value in its range. This contradicts $X(t)$'s being a function.

This complication is resolved by using a different parameter set. The parameter set used is the lexicographically ordered Cartesian product of the nonnegative reals and the natural numbers. The parameters are assigned in the following way. Consider an arbitrary sample path. At any time t in the nonnegative reals, if there are $n(n \geq 0)$ instantaneous state changes, then $X(t, 0)$ is the state before the first (if any) state change, $X(t, 1)$ is the state after the first state change, ... $X(t, n - 1)$ is the state after the $n - 1$th state change, $X(t, m), m \geq n$ is the state after the nth instantaneous state change. Since at most a countably infinite number of instantaneous state changes can occur, the process is never in two or more states at the same "time". Note that the parameter set of the embedded Markov Chain need only be the nonnegative integers with the nth parameter meaning the nth state change.

5. Analysis of The Dining Philosophers

The Dining Philosopher model is a well-known example which violates net restrictions in previous TPN models. Although performance of this system is largely hypothetical in nature, it serves to illustrate the capabilities of the GTPN analyzer, and it yields some insight into the timing behavior of the dining philosopher protocol.

A GTPN model of the 5 Dining Philosophers [13] is shown in Figure 5.1. The initial marking of the net shows all 5 philosophers thinking. We have analyzed the model with deterministic think times, τ_i, as shown, for each philosopher i. We have also used a slightly modified model (see Figure 3.1), to represent think times that are geometrically distributed with mean τ_i. After thinking, the philosopher competes for two forks which are shared with neighboring philosophers on the left and right, respectively. After acquiring the forks, the philosopher spends a deterministic amount of time eating (d_i). This cycle is repeated as many times as necessary to finish the meal. The firing frequencies, f_i, associated with transitions that model fork acquisition are used to compute the probabilities that various maximal sets of competing philosophers get the forks they require. These probabilities are calculated as described in Section 3.2. Note that all of these transitions are in the same Generalized Conflict Set. Also note that, to simplify the diagram, we have omitted the arcs which model releasing the forks at the end of the dining transitions.

20

Figure 5.1. Dining Philosophers GTPN Model

A unique resource is associated with each thinking and dining transition. The GTPN analyzer will compute the long run fractions of time ($U_{i,tan}$ and $U_{i,d}$) that philosopher i spends thinking and dining. From these measures, we can calculate the long run fraction of time philosopher i waits for forks ($U_{i,idl} = 1 - U_{i,tan} - U_{i,d}$).

Performance of the dining philosophers is maximized when time spent waiting for forks and expected time to complete the dinner are minimized. Note that time spent waiting for forks can be minimized if a subset of the philosophers agree not to eat until the other philosophers are finished. We consider this to be a less-than-ideal situation, even if the expected time to complete the meal is minimized, as compared with a situation where all philosophers dine and think together. Thus, for simplicity, we will assume the dinner ends at time *Dmin*, the expected time that the first philosopher(s) will complete the meal, and assess a penalty for philosophers who have not finished. We assume that each philosopher requires R units of time dining to complete the meal. Let $U_{max,d}$ be the maximum value of $U_{i,d}$ over all i. The time to complete the dinner is then: $Dmin = \frac{R}{U_{max,d}}$.

The Dining Philosopher model can be analyzed quickly for various think times, dining times, and firing frequencies (i.e. relative aggressiveness in grabbing forks). It is easy to assign deterministic delays which yield multiple recurrent classes. For example, if all philosoophers have $d_i = 3$, two non-neighboring philosophers have a deterministic think time of 3, and the remaining philosophers have a deterministic think time of 5, then the Markov Chain has four recurrent classes. However, each set of parameters we report for the model with deterministic think times yields a Markov Chain that has one recurrent class. Analysis of each Markov Chain thus yields a single value for each performance estimate.

We first consider the case of two classes of philosophers and deterministic think times. The first class of philosophers, formed by any two non-neighbors, thinks for N units of time and dines for N units of time. The second class of philosophers also dines for N units of time. In one experiment, we let N=3, and vary the duration of the think time of the second class of philosophers between 1 and 12 units of time. The firing frequencies, f_i for both classes of philosophers in this experiment are set equal to one. Figure 5.2 shows $U_{i,\tau}$, $U_{i,idle}$, and $U_{i,d}$ for the model. Our first observation is that the fractions of time the philosophers spend thinking, waiting, and dining, vary in a complex way with the input parameters. This behavior would be further emphasized if we included data points at each integer for $1 \leq \tau_{i2} \leq 12$. Reasoning about the behavior of the system for one parameter setting (i.e. when $\tau_{i2} = 3$), shows that after 9 units of time, the system reaches "steady state", in which two philosophers are thinking, two are dining, and one is waiting, (interchangeably), forever after. We note that this behavior is highly dependent on the relative delays in the model.

We investigated the complex behavior of system performance as a function of varying think times for the two classes of philosophers further. Figure 5.3 shows the total time spent thinking, idle, and dining, and the amount of time needed to complete any unfinished portions of the meal, for a few interesting parameter settings, assuming the time required to complete the meal, R=60. Starting with a "baseline" model ($\tau_i = d_i = N$), in 5.3(a), we see that Dmin=2.5 hours, of which each class of philosopher spends 60 min.(40%) eating, 60 min. thinking, and 30 min. waiting for forks. This corresponds to $\tau_{i2} = 3$ in Fig 5.2. All philosophers finish the meal at the same time.

22
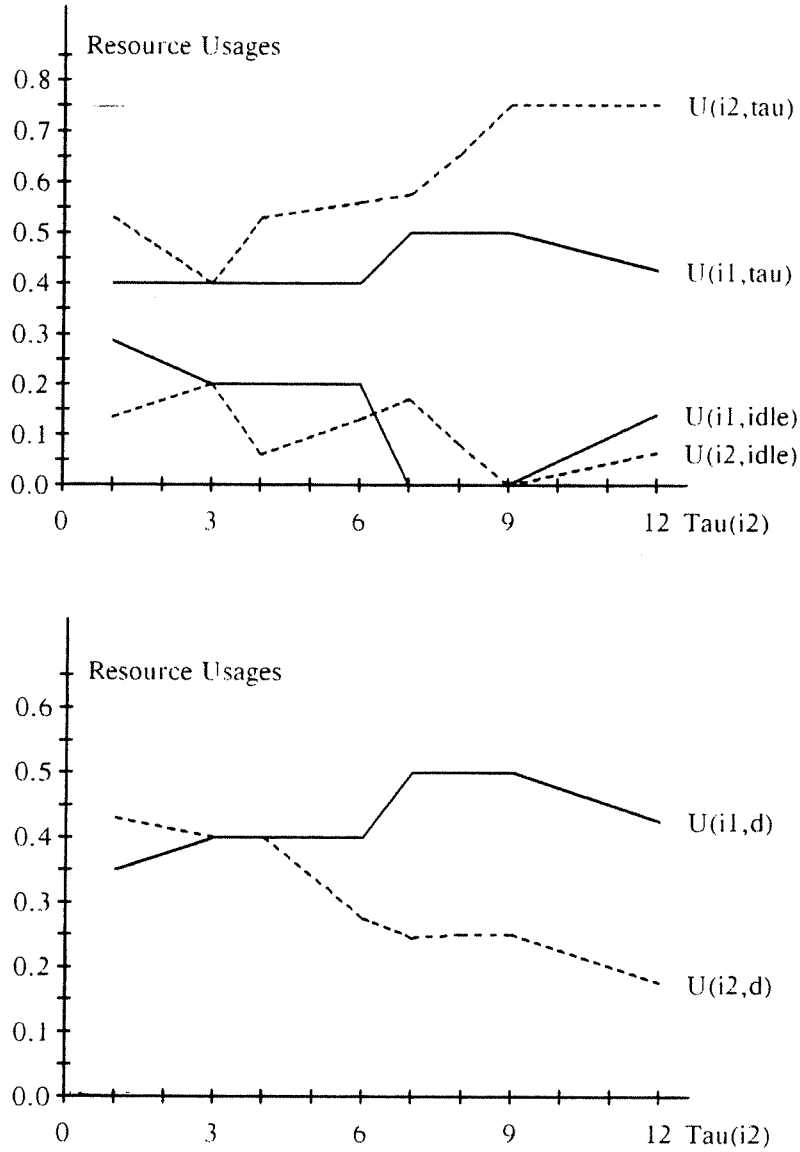
Figure 5.2. Performance Measures for varying deterministic think times. Two Classes of Philosophers. $\tau_{i1} = 3$ $\forall j$ $f_j = 1, d_j = 3$.

Note that two philosophers is the maximum number that can be dining at the same time, so the baseline model is optimum with respect to $U_{i,d} = 0.4$. The question is whether the idle time for the philosophers can be reduced while still dining at full capacity. In 5.3(b), the second class of philosophers reduce their idle time by slightly increasing their think time by some amount $x, x < \frac{N}{2}$. In 5.3(c) the second class of three philosophers increase their think times to 3N, which reduces idle time to zero for all philosophers, but causes the three to miss half their meal. This corresponds to $\tau_{i2} = 9$ in Figure 5.2. In 5.3(d), both classes of philosophers have think times set to $\tau_i = 1.5d_i = 1.5N$,
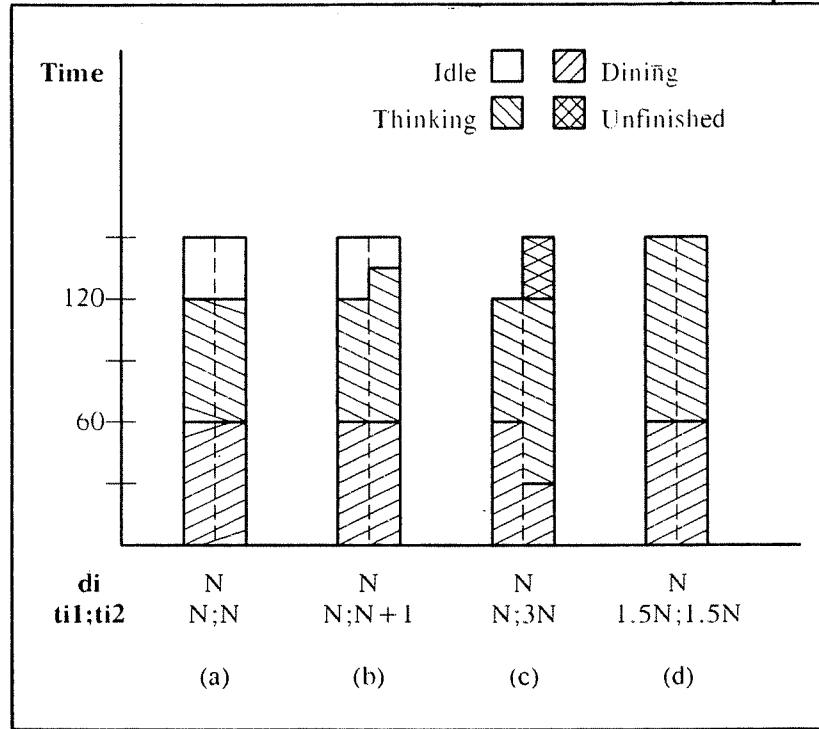
**Figure 5.3.** Performance for Selected Parameters
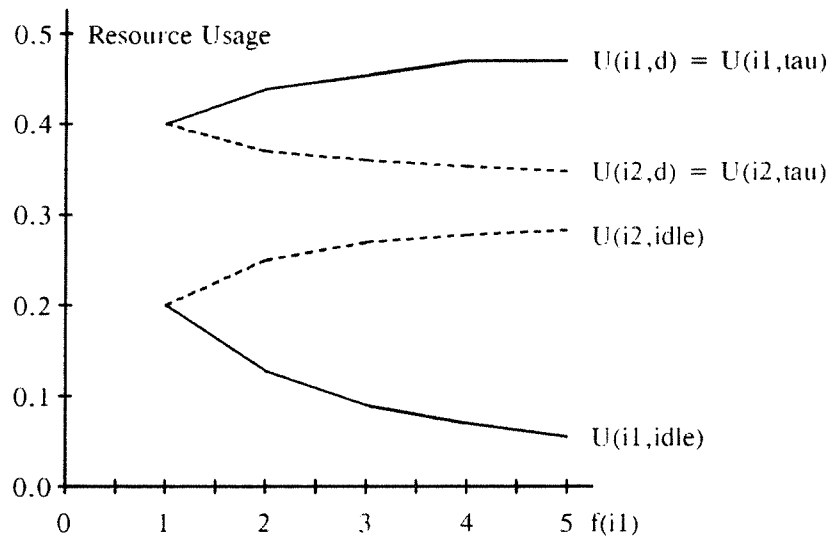
Figure 5.4. Performance with Aggressive Philosophers.

which represents the optimum behavior.

For the experiments above, we varied the think time while holding all dining times and firing frequencies constant. In the next experiment, we set the think time of the second class of philoso-

24

phers to N, and vary the firing frequencies of the first class of philosophers. Figure 5.4 gives the results of these analyses. Note that the maximum possible value for U_{il} for the given parameter settings, is 50%. When the firing frequency for the aggressive philosophers is 5.0 the fraction of time they spend waiting for forks is reduced by 70% (to 0.06).
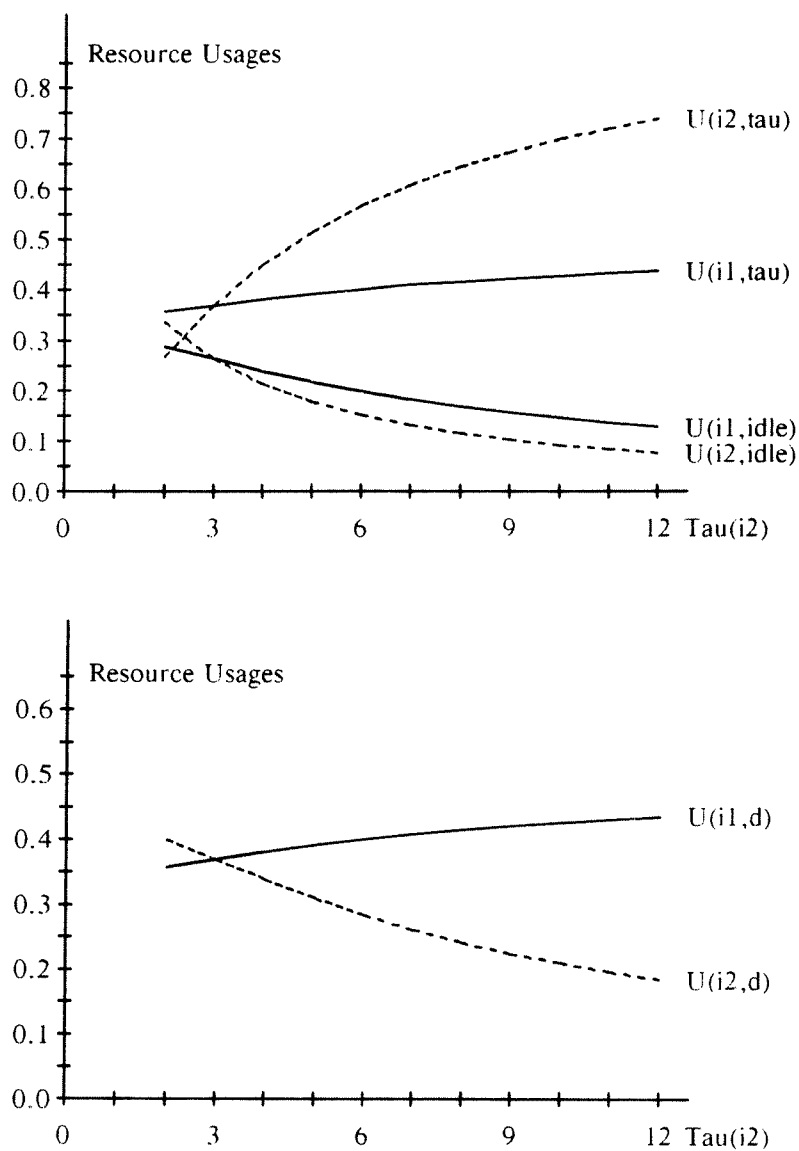
Figure 5.5. Varying geometric think times. Two Classes of Philosophers. $\tau_{i1} = 3$ $\forall j$ $f_j = 1, d_j = 3$.

Finally, we repeated the first experiment (Figure 5.2) with geometric think times. Figure 5.5 shows the performance estimates as a function of mean think time of the class two philosophers. The trends in the performance estimates as τ_{i2} varies are qualitatively the same as in the deterministic

25

model. However the performance curves are smooth, in contrast to the erratic variations in Figure 5.2. The erratic performance in the deterministic models is due to cyclic dependencies. These dependencies also make multiple recurrent classes more likely in the deterministic models.

6. Comparison with SPN Models

Models in which a transition's firing time is an exponential random variable, Stochastic Petri Nets (SPN), have been independently proposed by Natkin [5], Symons [6], and Molloy [7,8]. Marsan, Balbo, and Conte [10] generalized the continuous-time SPN model, GSPN, by allowing transitions which fire in zero time. Molloy [7,9] also proposed a discrete time SPN model with transition firing times that are geometric random variables. The SPN models are interesting because the reachability graph for these models are (continuous-time or discrete-time) Markov Chains. In this section, we compare the conflict resolution and probability assignment methods of the GTPN model and the SPN models. We then compare the modeling power of the GTPN and SPN models in four respects. Finally, we comment on the complexity issues concerning deterministic firing durations.

6.1. Conflict Resolution

The addition of timing information to the Petri net model provides several options for conflict resolution. The method of resolving conflicts in the GTPN is different than the method defined for the SPN models.

The GTPN conflict resolution submodel uses (possibly state-dependent) transition firing frequencies to resolve conflicts. Our underlying assumption is that the conflict must be resolved before one of the conflicting transitions can start firing. We also assume that once a transition is in progress, it cannot be preempted by a new conflict. This submodel is useful, for example, if the conflict is due to contention for a shared resource, or represents alternative outcomes based on resource availability.

In contrast, the conflict resolution submodel implicit in the SPN is based on competing transition delays. The transition which fires first wins the conflict. This mechanism is based on the view that the physical events modeled by the conflicting transition are in progress simultaneously and

that the completion of one event disables the other. For example, this view holds if one transition models the successful acknowledgement of a message in a computer network, and the other transition models a timeout process. The removal of tokens at the time a transition is enabled is not useful for this approach.

Each of the conflict resolution submodels is valid for the corresponding view of the physical system. Either can be employed in the SPN or the TPN model. Marsan's GSPN model uses firing frequencies for instantaneous transitions and competing delays for exponentially distributed transition rates. Razouk and Phelps use firing frequencies and *enabling times*. This is equivalent to implementing the two conflict resolution policies within one model. If it becomes useful, the GTPN reachability graph construction methods can easily be extended to incorporate enabling times.

6.2. Probability Assignment

The continuous time GSPN model and the GTPN have an embedded discrete time Markov Chain, while the discrete time SPN model itself is a discrete time Markov Chain. In either case, probabilities need to be assigned to next states. Each model assigns probabilities that are consistent with its conflict resolution semantics.

When the time until the next state is positive (i.e. no instantaneous transitions are firing), the SPN models assign probabilities according to competing transition delays. In the continuous time SPN's, with probability one, no two transitions in progress will finish firing simultaneously. This simplifies probability assignments. In particular, the probability of the next state associated with transition t_i's finishing first has probability equal to t_i's firing rate divided by the sum of the firing rates of the transitions that are in progress.

In the discrete time SPN, with probability greater than zero, two or more transitions in progress can finish firing simultaneously. Though more complicated, it is still possible to assign probabilities to next states using competing transition delays [7,9]. Unfortunately, using competing delays to assign next state probabilities restricts the allowed probability assignments in the important special case of deterministic firing delays. A deterministic firing delay is represented as a geometric firing

27

delay that has probability one of firing in the next time step. Consequently, using competing delays implies that all the next states have equal probability.

The GTPN uses *transition frequencies* to assign next-state probabilities, independent of transition delays. The transition frequency method, unlike the discrete time SPN, can assign non-uniform probabilities to next states in the case of conflicting transitions with deterministic delays. Competing deterministic delays lead to only one possible next state. Thus, assigning next state probabilities is trivial in this case (as long as we support enabling times). It is also possible to assign frequencies for conflicting geometric transitions such that assigned next-state probabilities are equal to probabilities based on competing delays. This is straightforward, for example, if conflict only arises at the first time step that both holding times are enabled. We can resolve the conflict with instantaneous transitions having the appropriate frequencies and then model the geometric delays in the usual way. In more complex cases the assignment becomes less straightforward. In general, although the frequency method is powerful, the assignment of static (state-dependent) frequencies, since they will be used for the dynamic calculation of probabilities, requires careful thought during model construction.

The *random switches* used in the GSPN when there are instantaneous transitions firing is a similar method which uses transition delays to determine probability assignments. In the random switches method, however, instead of one frequency expression per transition, a probability distribution is explicitly given for each possible set of enabled transitions.

6.3. Modeling Power

The GTPN has more modeling power than the existing SPN models in the probability assignments discussed above and in three additional respects: 1) firing durations, 2) resources, and 3) analysis of multiple recurrent classes.

The first respect is firing durations. The GTPN can represent deterministic firing durations which are arbitrary nonnegative real values, including zero. The GTPN can also model geometric holding times as can the discrete time SPN. Since a hold time in the discrete time SPN must be a

multiple of some unit step, the GTPN can be said to have more modeling power, with respect to firing durations, than the discrete time SPN. Also, since the geometric distribution is the discrete-time analog of the exponential distribution, and since the GTPN cannot represent deterministic delays, the GTPN can also be said to have more modeling power, with respect to firing durations, than the GSPN model.

The second respect is that the GTPN *resources* provide a simple method of calculating a wide range of performance estimates. The existing SPN models only provide steady state probabilities. It should, however, be fairly simple to add resources to an SPN model.

The third respect is the GTPN's analysis of multiple recurrent classes. The GTPN allows the performance evaluation of systems that have several possible long run behaviors. In contrast, the GSPN and the discrete time SPN assume that their Markov Chains are irreducible (i.e. have only one recurrent class). Again, we believe that the SPN analysis could be developed to support multiple recurrent classes.

We note that the Extended Stochastic Petri Net (ESPN) model of Dugan, Trivedi, Geist, and Nicola [11] is a SPN model that, in at least one respect, is more powerful than the GTPN. The ESPN allows arbitrary holding times and has been shown to be useful for analyzing system response to failure. However, the ESPN can only be analyzed if the reachability graph is acyclic, unless restrictions are placed on transitions which have delays that are not exponentially distributed.

6.4. Complexity Issues

Representing deterministic holding times inherently requires larger state spaces than when holding times are geometrically or exponentially distributed, because the memoryless property of the geometric and exponential random variables does not apply. Thus, both the GTPN and the discrete-time SPN model have the potential for large state spaces when deterministic holding times are represented.

The GTPN contains new states for start firing as well as end firing events. If we loosely identify these two state changes as one, there is an equivalence between the states in the GTPN and the

29

deterministic SPN when geometric holding times are in progress (i.e. one state change per time step, including a cycle back to a given state with the probability that none of the geometric delays completes in the step). When only deterministic holding times are in progress, however, the GTPN may not contain state changes (or new states) for every time step, whereas the deterministic-time SPN must. Thus, it appears that the GTPN has at most twice as large a state space as the discrete-time SPN and that for some deterministic models, the GTPN has a smaller state space. We note that the RFT vector, which allows a potential reduction in the size of the state space, adds minimal complexity. It is easy to assign new values to the RFT vector and to find the smallest value in it.

Due to the inherent complexity of deterministic delays, two important goals during our development of the GTPN were to minimize the size of the state space and to minimize the cost of constructing and analyzing it. These goals are reflected in the algorithms and performance analysis techniques presented in Sections 3 and 4. The GTPN state space is reduced by generating next states for *maximal* set of events that occur simultaneously (including multiple start-firings of a single transition). Another example of reducing the cost of building the state space is our definition of Generalized Conflict Sets (GCS) and the Partition algorithm. Generalized conflicts sets are state independent and thus need only be calculated once. We thus are able to avoid a large fraction of the conflict determination cost when we are calculating the next states of the reachability graph. In contrast, the definition used in the discrete time SPN is based on partitioning enabled transitions into sets that are *in conflict*. This partition is state dependent and thus must be computed for each state.

That we achieved our goal can be seen by example. We have been able to apply the GTPN to obtaining exact performance estimates of multiprocessor memory and bus interference [3]. Previously, researchers have viewed obtaining such exact estimates to be computationally intractable. In contrast, with the GTPN a multiprocessor model with 12 processors, 10 memories, 2 buses, and a geometric time between memory requests with mean of 5 time units has 2026 states and requires 274 seconds to build the reachability graph and analyze it for performance estimates.

7. Conclusions

We have developed a generalized Timed Petri Net model for studying the performance of computer systems. The model imposes no restrictions on the net except that the state space be finite. We have also defined methods for performance analysis of the GTPN based on analysis of the embedded discrete-time Markov Chain. The analysis is defined for chains which include multiple recurrent classes. Our comparison of this model with the discrete-time SPN and the continuous-time GSPN models shows that the GTPN has more modeling power. In particular, we showed that the GTPN is a better bridge between the TPN and SPN models than is the discrete-time SPN model. We have implemented the GTPN Analyzer and demonstrated that the algorithms presented are efficient.

We expect that the capability to obtain accurate analytic performance estimates will yield insight into many computer system models. Such exact estimates have already been obtained with the GTPN for multiprocessor memory and bus interference [3]. Future research plans include studying the applicability of GTPNs to other issues in computer performance evaluation, such as the performance of network protocols, the performance of load balancing algorithms in distributed systems, the performance of database machines, and the performance of advanced architectures for high-speed numeric or symbolic computation.

Acknowledgements

Many people have made helpful comments and suggestions during our research. We would like to especially thank Josh Chover, Bill Kalsow, Tom Kurtz, and Udi Manber.

References

[1] W.M. Zuberek, "Timed Petri nets and preliminary performance evaluation," in *Proc. 7th Ann. Symp. Comp. Arch.*, pp. 88-96, 1980.

[2] R.R. Razouk and C.V. Phelps, "Performance Analysis Using Timed Petri Nets," in *Proc. 1984 Int Conf. on Parallel Processing*, pp. 126-129, August, 1984.

[3] M.A. Holliday and M.K. Vernon, *Exact Performance Estimates for Multiprocessor Memory and*

Bus Interference, Tech.Rep. 593, CS Dept., UW-Madison, May 1985.

[4] C.A. Petri, "Kommunikation mit Automaten," *Schriften des Rheinisch-Westfalischen Institute fur Instrumentelle Mathematick an der Universitat Bonn,* Heft 2, Bonn, W. Germany, 1962; translation: C.F. Greene, Supplement 1 to Tech. report RADC-TR-65-337, Vol. 1, Rome Air Development Center, Grifiss Air Force Base, NY 1965.

[5] S. Natkin, "Reseaux de Petri stochastiques", these de Docteur-Ingenieur, CNAM-Paris, June 1980.

[6] F.J.W. Symons, "Introduction to numerical Petri nets, a general graphical model of concurrent processing systems," A.T.R., vol. 14, Jan. 1980.

[7] M.K. Molloy, "On the integration of delay and throughput measures in distributed processing models," Ph.D dissertation, Univ. California, Los Angeles, 1981.

[8] M.K. Molloy, "Perf. Analysis using stochastic Petri nets," *IEEE Trans. Comp.,* vol. C-31, pp. 913-917, Sept. 1982.

[9] M.K. Molloy, "Discrete Time Stochastic Petri Nets," *IEEE Trans. Soft. Engr.,* vol. SE-11, pp. 417-423, April 1985.

[10] M.A. Marsan, G. Balbo, and G. Conte, "A Class of Generalized Stochastic Petri Nets," *ACM Trans. on Computer Systems,* vol. 2, pp. 93-122, May 1984.

[11] J.B. Dugan, K.S. Trivedi, R.M. Geist, and V.F. Nicola, "Extended Stochastic Petri Nets: Applications and Analysis," *Performance 84,* pp. 507-519, Paris, France, December 1984.

[12] M.A. Holliday and M.K. Vernon, "A Generalized Timed Petri Net Model for Performance Analysis," in *Proc. Int. Workshop on Timed Petri Nets,* July 1985.

[13] J.L Peterson, *Petri Net Theory and the Modeling of Systems.* Englewood Cliffs, NJ: Prentice-Hall, 1981.

[14] C. Ramachandani, "Analysis of Asynchronous Concurrent Systems by Timed Petri Nets," Ph.D. Dissertation, MIT, 1974.

[15] C.V. Ramamoorthy and G.S. Ho. "Performance Evaluation of Asnchronous Systems using Petri nets," *IEEE Trans. Software Engr.*, vol. SE-6, pp. 440-449, Sept. 1980.

[16] G.J. Nutt, "Evaluation nets for computer systems perf. analysis," in 1972 *Fall Joint Computer Conf., AFIPS Conf. Proc.*, vol.41., Montvale, N.J.: AFIPS Press, 1972, pp. 279-286.

[17] J.G. Kemeny and J.L. Snell, *Finite Markov Chains.* New York, NY: Springer-Verlag, 1976.

[18] R. Sedgewick, *Algorithms.* pp. 428-430, Reading, MA: Addison-Wesley, 1983.

[19] H.M. Taylor and S. Karlin, *An Introduction to Stochastic Modeling.* Orlando, FL: Academic Press, 1984.

[20] C.H. Sauer and K.M. Chandy, *Computer Systems Perf. Modeling*, Englewood Cliffs, NJ: Prentice-Hall, 1981.