COMPUTER SCIENCES
DEPARTMENT

University of Wisconsin -
Madison

PROPERTIES OF CONFLICT FREE AND PERSISTENT
PETRI NETS

by

L. H. Landweber
and
E. L. Robertson

Computer Sciences Technical Report #264

December 1975

PROPERTIES OF CONFLICT FREE AND PERSISTENT PETRI NETS*

by

L. H. Landweber
Computer Sciences Department
University of Wisconsin--Madison

and

E. L. Robertson
Computer Science Department
Pennsylvania State University

*This research was partially supported by NSF Contract #GJ33087.

Abstract

Petri nets have been extensively studied because of their suitability as models for asynchronous computing. Despite this effort, the mathematical properties of Petri nets are not very well understood. In this paper, we investigate two important special types of Petri nets, the conflict free and the persistent nets, the former being a proper subset of the latter. Our results completely characterize the sets of reachable markings attainable by such nets. Reachability sets of persistent nets are shown to be semi-linear. A stronger result is obtained for conflict free nets which results in an exponential time algorithm for deciding boundedness of such nets. The best known upper bound for deciding boundedness of arbitrary nets is Ackermann's function. We conclude with a proof that all reachability sets of Petri nets may be realized with a minimal amount of non-persistence.

1. Introduction

Petri nets and vector addition systems have been extensively studied because of their suitability as models of asynchronous computing. Despite this effort, the mathematical properties of these equivalent models are not very well understood. For example, it is not known whether the reachability and liveness problems are decidable although the recursive equivalence of these problems has been demonstrated [2]. It is known [5] that if reachability is decidable, then the time required by the decision procedure must be at least exponential in the size of the inputs (i.e., the size of a representation of the Petri net plus the marking to be reached).

In this paper, we investigate two important special types of Petri nets, the conflict free and the persistent nets, the former being a proper subset of the latter. Our results completely characterize the sets of reachable markings attainable by such nets. Persistence for Petri nets is very similar to the persistence property used by Lipton, Miller and Snyder [6] in their study of linear asynchronous structures. It is also similar to a notion used by Muller to study various types of switching circuits.

Liveness is known to be decidable for persistent nets [4] while both liveness and reachability are decidable for conflict free nets [1]. Our first main result (Section 4) shows that persistent nets have semi-linear sets of reachable markings. Since conflict free nets are persistent, this answers a question left open in [1]. Our proof does not yield a decision procedure for reachability in persistent nets because the construction is not effective. However, it does indicate a strong probability that such a decision procedure does exist. Semi-linearity has played an important role in the study of Petri nets. For example, Van Leeuwen's [9] decision procedure for the reachability problem for 3 coordinate vector addition systems involves the construction of semi-linear representations for sets of reachable points. Rabin's proof that containment of sets of reachable points for arbitrary Petri nets is undecidable involves building nets which have non semi-linear reachability sets. We believe that the relationship we exhibit between persistence and semi-linearity, plus the fact that only a "small" amount of non-persistence is necessary to achieve any (non-semi-linear) reachable set of markings (Section 6), clearly indicates why it has been so difficult to obtain significant results for arbitrary Petri nets.

We next (Section 5) investigate properties of conflict free nets. In particular, we show that for any conflict free net, there is a constant c such that for an arbitrary initial marking with x tokens, any place in the net receives either at most cx tokens or an unbounded number of tokens. A corollary to this result is an exponential time algorithm for deciding boundedness for conflict free nets. This may be contrasted with Karp and Miller's [3] algorithm for deciding boundedness of arbitrary nets which has Ackermann's function as a time complexity bound.

In Section 2 we give some basic definitions and notation. Section 3 contains some combinatorial lemmas for various classes of Petri nets. Sections 4 and 5 deal with persistent and conflict free nets respectively. In the last section, we discuss some open problems and give examples which illustrate the central role that persistence and semi-linearity play in the study of Petri nets.

2. Definitions and Notation

A Petri net is a quadruple.

$$P = \langle P,T,A,M_0 \rangle$$

where P is a finite set of places; T is a finite set of transitions or firing bars; A is a finite set of arcs, $A \subseteq (P\times T) \cup (T\times P)$; and $M_0: P \to N$, N the set of natural numbers, is the initial marking.

Initially each place p of the Petri net contains $M_0(p)$ tokens. Let t be a transition. Then $\{p|(p,t)\epsilon A\}$ and $\{p|(t,p)\epsilon A\}$ are called the input places (inputs) and output places (outputs) respectively of t.

Transition t is enabled or fireable when each input place of t contains at least one token. If t is enabled, then it may be fired which results in the removal of one token from each input place of t and the addition of one token to each output place of t. If t is not enabled, then it is disabled. Write $M_1 \xrightarrow{t} M_2$ ($M_1 \xrightarrow{t}$) to indicate that t is enabled by the marking M_1 and that the firing of t yields marking M_2 (t is enabled by the marking M_1). Extend the notation and definitions to sequences of transitions, $\sigma \in T^*$, called firing sequences.

The set of reachable markings or the reachability set R_P of the Petri net $P = \langle P,T,A,M_0 \rangle$ is $\{M|M_0 \xrightarrow{\sigma} M$, for some $\sigma \in T^*\}$. If $M \in R_P$ we say that M is reachable in P. The reachability problem for a class C of Petri nets is the problem of deciding, given an arbitrary $P \in C$ and marking M, whether $M \in R_P$.

A place in a Petri net is bounded if there is a $c \in N$ such that for all reachable markings M, $M(p) \le c$. A Petri net is bounded if each place in the net is bounded. A Petri net (place in a Petri net) is unbounded if the net (the place) is not bounded. A subset P_1 of the set of places is simultaneously unbounded if for each $n \in N$ there is a reachable marking M such that $M(p) \ge n$ for each $p \in P_1$.

A Petri net is persistent if for all $t_1,t_2 \in T$, $t_1 \ne t_2$ and any reachable marking M, $M \xrightarrow{t_1}$ and $M \xrightarrow{t_2}$ imply $M \xrightarrow{t_1 t_2}$. I.e., if t_1 and t_2 are enabled at a reachable marking, then the firing of one cannot disable the other.

A place p and a transition t are or a self loop if p is both an input place and an output place of t.

A Petri net is conflict free if every place which is an input of more than one transition is on a self loop with each such transition. Conflict free nets are persistent though the converse need not be true.

A set M of markings is linear if there is a finite set of functions $\{f_i|f_i:P \to N\ (0\le i\le n)\}$ such that

$$M = \{f_0 + \sum_{i=1}^{n} c_i f_i | c_i \ge 0,\ 1 \le i \le n\}.$$

M is semi-linear if it is a union of a finite number of linear sets.

3. Combinatorial Properties

In this section we obtain some combinatorial properties of persistent and arbitrary Petri nets. In the following, assume that $P = <P,T,A_0,M_0>$ is a fixed but arbitrary Petri net with transitions $T = \{t_1,\ldots,t_k\}$. Define the Parikh map (see [7], for the first use of this important idea) $PK: T^* \to N^k$ so that $PK(\sigma)_i$ is the number of occurrences of t_i in σ. The corresponding Parikh space T of P is

$$\{PK(\sigma) | \sigma \in T^*, \sigma \text{ is fireable}\}.$$

Observe that we have discarded certain information with this map, namely the sequence in which the transitions fire. We first show that the Parikh space of a persistent net is a lattice with respect to the partial order defined by: $x \le y$ $(x,y \in N^k)$ if all coordinates of x are less than or equal to the corresponding coordinates of y. Also $x < y$ if $x \le y$ and $x \ne y$. In the following all arithmetic operations on vectors are to be interpreted as being performed independently on all corresponding coordinates of the vectors used.

For σ, $\tau \in T^*$, $\sigma = a_1 \ldots a_n$ define $(\tau \dot- \sigma)$ as follows: Let τ_0 be τ. Obtain τ_{i+1} by deleting the leftmost occurrence of a_i from τ_i, if a_i occurs in τ_i. If not, then $\tau_{i+1} = \tau_i$. Define $(\tau \dot- \sigma)$ to be τ_n.

Following an argument similar to that used by Keller [4] we have:

Lemma 3.1. Let σ and τ be fireable sequences in a persistent net. Then there is a fireable sequence β such that

$$PK(\beta) = \underline{\max} \; (PK(\sigma), PK(\tau)).$$

Moreover, β may be constructed so that $\beta = \sigma \cdot (\tau \dot- \sigma)$. □

The following trivial example (Figure 1) shows that the same result does not hold for min.

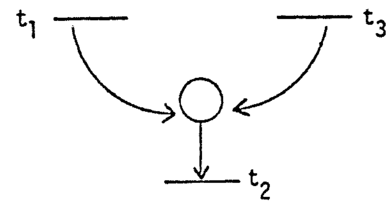

Figure 1

Clearly both t_1t_2 and t_3t_2 are fireable sequences with corresponding Parikh space points (1,1,0) and (0,1,1). But the transition t_2 corresponding to (0,1,0) is not itself fireable.

Theorem 3.2. For a persistent Petri net, the corresponding Parikh space is a lattice under the natural ordering of vectors ($x \le y$ iff $x_i \le y_i$ for all i, $1 \le i \le k$).

Proof. Join in the lattice is simply max by Lemma 3.1. Meet is not necessarily min, by the above remark, but can be shown unique. Consider two points x,y of the Parikh space of the net. If two points u,v of the Parikh space satisfy $u \le x$, $u \le y$, $v \le x$, $v \le y$, then $w = \max(u,v)$, satisfies $w \le x$, $w \le y$ so w is also a lower bound of x,y. But by Lemma 3.1, w is in the Parikh space of the net. □

Figure 2 shows an example of a non-persistent net whose Parikh space is {(0,0),(1,0),(0,1)}. Since (1,0) and (0,1) do not have a join, the space is not a lattice.

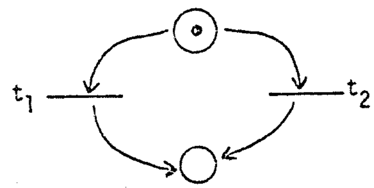

Figure 2

The following lemmas will be useful in Sections 4 and 5. A sequence $\alpha = t_1 t_2 \ldots \in T^\omega$ is an ω-firing sequence if every finite prefix of α is a fireable firing sequence. If there exists an ω-firing sequence α such that arbitrarily large markings of p (of every member of $P_1 \subseteq P$) may be obtained by firing finite prefixes of α, then we say that p (P_1) is <u>unbounded</u> (<u>simultaneously unbounded</u>) <u>on α</u>. The reader should compare this with the definitions of "unbounded place" and "simultaneously unbounded set of places" given in Section 2.

<u>Notation:</u> α_i is the i-th element of $\alpha = \alpha_1 \alpha_2 \ldots$.

$\alpha|i = \alpha_1 \ldots \alpha_i$ and $\alpha|ij = \alpha_i \ldots \alpha_j$.

It is possible to give a weaker definition of simultaneously unbounded on α, corresponding to the case where finite prefixes of an ω-firing sequence α mark places p_1 and p_2 with aribtrarily large values, but whenever the marking of P_1 is large, then the marking of P_2 is small and vice versa. The following lemma shows that the two notions coincide for arbitrary nets.

<u>Lemma 3.3</u>. Let A and B be sets of places of an arbitrary Petri net. Let A and B each be simultaneously unbounded on some ω-firing sequence α. Then there is an ω-firing sequence $\bar{\alpha}$ such that $A \cup B$ is simultaneously unbounded on $\bar{\alpha}$.

<u>Proof</u>. The idea is to find repeatable finite subsequences in α which increase A's markings without decreasing B's marking and vice versa. Let M_i be the marking at time i, i.e., after $\alpha|i$ has fired. Since A and B are each simultaneously unbounded on α, there must be times $a_1 < b_1 < a_2 < b_2 < \ldots$ such that

a) $M_{a_{i+1}}(x) > M_{a_i}(x)$ for each $x \in A$

b) $M_{b_{i+1}}(x) > M_{b_i}(x)$ for each $x \in B$.

Let $p_1, \ldots, p_r$ be an enumeration of $P - A$ and $q_1, \ldots, q_s$ an enumeration of $P - B$. Choose an infinite subsequence $\{a_i'\}$ of the a_i's such that $M_{a_{i+1}'}(p_1) \geq M_{a_i'}(p_1)$ for $i \geq 1$. Similarly choose

an infinite subsequence $\{b_i'\}$ of the b_i's such that $M_{b_{i+1}'}(q_1) \geq M_{b_i'}(q_1)$ for $i \geq 1$. Iterate this procedure for $p_2, \ldots, p_r, q_2, \ldots, q_s$. This yields infinite subsequences $\{\bar{a}_i\}$ and $\{\bar{b}_i\}$ of $\{a_i\}$ and $\{b_i\}$ respectively such that $M_{\bar{a}_{i+1}} \geq M_{\bar{a}_i}$ and $M_{\bar{b}_{i+1}} \geq M_{\bar{b}_i}$ for $i \geq 1$. From $\{\bar{a}_i\}$, $\{\bar{b}_i\}$ choose times $s_i < t_i < u_i < v_i$, $v_i < s_{i+1}$ for $i \geq 1$ which satisfy:

a) $M_{t_i}(x) > M_{s_i}(x)$ for each $x \in A$

$M_{t_i}(x) \geq M_{s_i}(x)$ for each $x \in P$

b) $M_{v_i}(x) > M_{u_i}(x)$ for each $x \in B$

$M_{v_i}(x) \geq M_{u_i}(x)$ for each $x \in P$.

Let $\beta_i = (\alpha|(s_i+1)t_i)$

$\gamma_i = (\alpha|(t_i+1)u_i)$

$\delta_i = (\alpha|(u_i+1)v_i)$

$\epsilon_i = (\alpha|(v_i+1)s_{i+1})$.

Note that $\alpha = (\alpha|s_1)\ \beta_1\gamma_1\delta_1\epsilon_1\beta_2\gamma_2\delta_2\epsilon_2 \ldots$. Then the required $\bar{\alpha}$ is

$$\bar{\alpha} = (\alpha|s_1)\beta^2{}_1\gamma_1\delta^2{}_1\epsilon_1\beta^2{}_2\gamma_2\delta^2{}_2\epsilon_2 \ldots$$

. Each new $\beta_i(\delta_i)$ increases the marking of each place in A (each place in B) without decreasing the markings of places in B(places in A). At all times after ϵ_i has been fired, all places in $A \cup B$ will have at least i tokens. Hence $A \cup B$ is simultaneously unbounded on $\bar{\alpha}$. □

<u>Corollary 3.4</u>. Let A be a set of places which is simultaneously unbounded on some ω-firing sequence α. Then there is an ω-firing

sequence $\bar{\alpha}$ such that for all $n \in N$, there is an i_n, such that for $i > i_n$, $\bar{\alpha}|i$ marks each member of A with at least n tokens.

Proof. This follows immediately from the proof of Lemma 3.3. □

The next two results are true for persistent nets.

Theorem 3.5. Let P be a persistent net with $A \subseteq P$, $B \subseteq P$. Then the following are equivalent.

1. A and B are simultaneously unbounded on some ω-firing sequence.
2. A and B are simultaneously unbounded.
3. $A \cup B$ is simultaneously unbounded.
4. $A \cup B$ is simultaneously unbounded on some ω-firing sequence.

Proof. $2 \Rightarrow 1$: Assume A and B are simultaneously unbounded. We construct an ω-firing sequence on which both A and B are simultaneously unbounded. Let α_1 be a finite fireable firing sequence which marks each place in A with ≥ 1 token. Let β_1 be a finite fireable firing sequence which marks each B member with $\geq 1 + |\alpha_1|$ tokens where $|\alpha_1|$ is the length of α_1. By Lemma 3.1, because P is persistent, $\sigma_1 = \alpha_1 \cdot (\beta_1 \dot{-} \alpha_1)$ is fireable. But σ_1 marks each member of B with ≥ 1 token. Assume α_i, β_i, σ_i have been defined for $1 \leq i < n$ and satisfy:

a) α_i is fireable and marks each A member with $\geq i$ tokens.

b) β_i is fireable and marks each B member with $\geq i + |\alpha_i|$ tokens.

c) $\sigma_i = \alpha_i \cdot (\beta_i \dot{-} \alpha_i)$ is fireable and marks each B member with $\geq i$ tokens.

d) σ_i is an initial segment of σ_{i+1} for $1 \leq i < n - 1$.

Let γ be a fireable sequence which marks each A member with $\geq n + |\sigma_{n-1}|$ tokens. By Lemma 3.1, $\alpha_n = \sigma_{n-1} \cdot (\gamma \dot{-} \sigma_{n-1})$ is fireable and marks each member of A with $\geq n$ tokens so α_n satisfies a). Now let β_n be a fireable sequence which marks each B member with $\geq n + |\alpha_n|$ tokens (satisfying b)). By Lemma 3.1, $\sigma_n = \alpha_n \cdot (\beta_n \dot{-} \alpha_n)$ is fireable and marks each B member with $\geq n$ tokens (satisfying c)). Moreover σ_n and σ_{n-1} satisfy d). In the limit, we obtain an ω-firing sequence on which both A and B are simultaneously unbounded.

$3 \Rightarrow 2$: obvious from definitions.

$4 \Rightarrow 3$: obvious from definitions.

$1 \Rightarrow 4$: Lemma 3.3 □

Corollary 3.6. Let P be a persistent Petri net with $A \subseteq P$. Then A is simultaneously unbounded if and only if A is simultaneously unbounded on some ω-firing sequence.

Lemma 3.7. Let P be an arbitrary Petri net. A transition t occurs ω times in some ω-firing sequence α if and only if either:

a) t is on a loop all of whose transitions occur ω times in α (t is on an α-ω-loop); or

b) t is on a path from a loop such that all transitions on the loop and all transitions on the path up to and including t occur ω times in α (t depends on an α-ω-loop); or

c) t is on a path all of whose transitions, including t, occur ω times in α and the first transition on the path has no input places (t is on an α-ω-path).

Proof. $\Leftarrow$ trivial

$\Rightarrow$ Assume t occurs ω times in α. Each input place to t must be an output place for some transition which occurs ω times in α. Iterate the backward search. Eventually, either a transition repeats (cases a or b) or a transition is reached which has no input places (case c). □

Lemma 3.8. Let P be a persistent net with transitions $t_1, t_2 \in T$. The following are equivalent:

1. For each $k \geq 1$, there are finite fireable sequences β_1^k and β_2^k having at least k occurrences of t_1 and t_2 respectively.

2. For each $k \geq 1$ there is a finite fireable sequence β in which t_1 and t_2 each occur at least k times.

3. t_1 and t_2 occur infinitely often on some ω-firing sequence α.

4. There are ω-firing sequences α_1 and α_2 such that t_1 and t_2 occur infinitely often on α_1 and α_2 respectively.

Proof. The proof is similar to proofs given in [4]. □

Lemma 3.9. Let $\alpha t\delta t'$ be a fireable sequence of a conflict free Petri net P where $t, t' \in T$ and $\delta \in T^*$. Then there is a fireable sequence $\alpha t\beta t'$ such that neither t nor t' occurs in β.

Proof. The proof is by induction on $L(t,\delta,t')$, the maximum length of directed paths in P from t to t' which are subpaths of δ. Let $L(t,\delta,t')$ be 0 if no directed path from t to t' is a subpath of δ.

If $L(t,\delta,t') = 0$, then no path from t to t' is a subpath of δ. First delete all occurrences of t in δ. Then iterate the process of deleting occurrences of other transitions in $\alpha t\delta$ which have become disabled. Call the portion of δ that remains β. The sequence $\alpha t\beta t'$ is fireable because, for every deleted transition $\bar{t}$, the firing of $\bar{t}$ could not have been necessary to the firing of t'. (Else some subpath of δ would connect t to t'.) Now let δ be the largest initial segment of β not containing t'.

Assume the lemma is true for all transitions $\bar{t}$ and all sequences γ such that $L(t,\gamma,\bar{t}) \leq n$ and let $L(t,\delta,t') = n+1$. Assume that t' does not occur in δ (if it does, then let δ be the longest initial segment of δ not containing t'). Let $t_1,\ldots,t_r$ be those transitions which occur in δ and which have output places that are also input places of t'. If there are no such transitions, then $\alpha tt'$ is fireable. Let δ_i be the initial segment of δ which precedes the last occurrence of t_i in δ. Then, for each i, $\alpha t\delta_i t_i$ is fireable and $L(t,\delta_i,t_i) \leq n$ so by the induction hypothesis, there is a β_i not containing t such that $\alpha t\beta_i t_i$ is fireable. But then by Lemma 3.1, there is a fireable sequence σ which begins with αt and which satisfies

$$PK(\sigma) = \max \{PK(\alpha t\beta_i t_i)\}.$$

In particular, σ is of the form $\alpha t\beta$ where β contains no occurrences of t or t' and at least one occurrence of each t_i. But the net is conflict free so $\alpha t\beta$ enables t' and hence $\alpha t\beta t'$ is fireable. The Lemma is therefore true for $L(t,\delta,t') = n + 1$ and so by induction it is true for all n. □

Lemma 3.10. Let P be a conflict free Petri net. Let t, t' be transitions which occur infinitely often in the ω-firing sequence α. Then there is an ω-firing sequence α' in which t and t' occur infinitely often where

$$\alpha' = \gamma t\beta_1 t'\gamma_1 t\beta_2 t'\gamma_2 t \ldots$$

and t does not occur in γ, β_i and γ_i $(i \geq 1)$.

Proof. The sequence can be written in the form

$$\alpha = \gamma t\delta t'\bar{\alpha}$$

where γ does not contain t and δ does not contain t'. The sequence α' is constructed by successive applications of Lemma 3.1 and 3.9.

Begin by applying Lemma 3.9 to obtain a fireable sequence $\gamma t\beta t'$ where β does not contain t. By Lemma 3.1, the ω-sequence σ given by

$$\begin{aligned}\sigma &= \gamma t\beta t' \cdot (\gamma t\delta t' \dot{-} \gamma t\beta t')\bar{\alpha} \\ &= \gamma t\beta t' \cdot (\delta \dot{-} \beta) \cdot \bar{\alpha} \\ &= \gamma t\beta t'\tau \cdot t\,\zeta\, t' \cdot \bar{\alpha}' \\ &= \gamma t\beta t'\tau \cdot t\,\zeta\, t' \cdot \bar{\alpha}'\end{aligned}$$

is fireable where τ does not contain t and ζ does not contain t'. Now repeat the above procedure using $t\zeta t'$ in σ. The iteration of this procedure yields the required sequence α'. □

4. Persistent Nets

In this section we show that the set of reachable markings of a persistent Petri net is semi-linear. In the following, let $P = \langle P,T,A,M_0 \rangle$ be a fixed but arbitrary persistent Petri net with k transitions and ℓ places. For convenience, we introduce an extended Parikh map, $EPK: T^+ \to N^{k+\ell}$. Let $\Sigma: T^+ \to N^\ell$ be defined by $\Sigma(\theta) = (m_1, \ldots, m_\ell)$, if θ is a fireable firing sequence and m_i is the change in the number of tokens in the place p_i as a result of firing θ. Then, for a fireable sequence θ,

$$EPK(\theta) = PK(\theta) \times \Sigma(\theta).$$

$EPK(\theta)$ is a $k + \ell$ - tuple whose coordinates give the number of occurrences of each transition in θ and the change in the marking of each place which results from firing θ. Recall that for $b,c \in N^{k+\ell}$, $b \le c$ if every coordinate of b is $\le$ the corresponding coordinate of c.

Definition. For the marking $v \in N^\ell$, let S_v be the set of points $b \in N^{k+\ell}$ for which there is a fireable $\theta \in T^+$ satisfying $EPK(\theta) = b$ and $\Sigma(\theta) \ge 0$.

Definition. For each marking $v \in N^\ell$, let F_v be the set of minimal points in S_v. I.e.,

1. $F_v \subseteq S_v$
2. if $b_1, b_2 \in F_v$, then b_1 and b_2 are incomparable ($b_1 \nleq b_2$, $b_2 \nleq b_1$)
3. for each $c \in S_v$, there is a $b \in F_v$, $b \le c$.

Lemma 4.1. Each F_v is finite.

Proof. By Koenig's infinity Lemma, every set of pairwise incomparable vectors on $N^{k+\ell}$ is finite. Hence each F_v is finite. □

Lemma 4.2. For $u,v \in N^\ell$, if $u \ge v$, then $(\forall a \in F_v)(\exists b \in F_u)[b \le a]$.

Proof. Let $a \in F_v$. Then $EPK(\theta) = a$ for some θ fireable at v. But θ is also fireable at u so either $a \in F_u$ or some $b < a$ is in F_u. □

Observe that F_u can only contain points that are smaller than or incomparable with the points in F_v. This is true because when the marking is increased from v to u, all previously enabled sequences remain enabled while some additional sequences may become enabled.

Lemma 4.3. $F = \bigcup_{v \in N^\ell} F_v$ is finite.

Proof. Let MF be the set of minimal points of F. Koenig's Infinity Lemma implies that MF is finite. Hence there is a finite $S \subseteq N$ such that

$$MF \subseteq \bigcup_{v \in S} F_v.$$

Let $w = \underline{\max}\ S$ (componentwise maximum). We first show that $MF = F_w$.

1. $MF \subseteq F_w$: Let $a \in MF$. Then $a \in F_v$ for some $v \in S$. Since $w \ge v$, by Lemma 4.2, there is a $b \in F_w$, $b \le a$. But $b < a$ is impossible because $a \in MF$. Therefore $a = b$ so $a \in F_w$.
2. $F_w \subseteq MF$: Let $a \in F_w$. Since MF is the minimal set for $\bigcup_{u \in N^\ell} F_u$, there is a $b \in MF$ such that $b \le a$. But then $b \in F_v$ for some $v \in S$. Since $w \ge v$ there is a $c \in F_w$, $c \le b$. Combining the above we get $c \le b \le a$ where $a,c \in F_w$, $b \in MF$. But F_w is a set of minimal points so $a = b = c$ and $a \in MF$.

Because $F_w = MF$, $u \ge w$ implies that $F_u = MF$. Hence

$$F = \bigcup_{v \in N^\ell} F_v = \bigcup_{v \le w} F_v$$

so F is finite by Lemma 4.1. □

The next lemma shows that the effect of any firing sequence can be achieved by concatenating canonical sequences whose images under EPK are in the finite set F.

Lemma 4.4. Let $\sigma \in T^+$ be fireable at some marking $v \in N^\ell$ and assume $\Sigma(\sigma) \geq 0$. Then there is a sequence $\xi_1 \ldots \xi_r$, $r \geq 1$, $\xi_i \in T^+$ which satisfies

1. $EPK(\xi_i) \in F \quad 1 \leq i \leq r$
2. $\Sigma(\xi_1 \ldots \xi_r) = \Sigma(\sigma)$
3. $\xi_1 \ldots \xi_r$ is fireable at v.

Call $\xi_1 \ldots \xi_r$ a decomposition of σ.

Proof. The proof is by induction on the length of σ.

$|\sigma| = 1$: $EPK(\sigma)$ is minimal so $\sigma = \xi_1$.

Assume the result holds for sequences of length $< n$ and let $|\sigma| = n$. If $EPK(\sigma) \in F_v$, then $EPK(\sigma) \in F$ and we are done. If $EPK(\sigma) \notin F_v$, then there is some η fireable at v such that $EPK(\eta) \in F_v$ and $EPK(\eta) < EPK(\sigma)$ (F_v is the set of minimal points in $N^{k+\ell}$ corresponding to sequences that are fireable at v and which do not decrease the marking.). Because the net is persistent, Lemma 3.1 implies that there is a $\delta \in T^+$ such that $\eta\delta$ is firable at v and $EPK(\sigma) = EPK(\eta\delta)$. Since $\eta\delta$ is fireable at v, δ is fireable at $v + \Sigma(\eta)$. Also $EPK(\eta) < EPK(\sigma) = EPK(\eta\delta)$ and $\Sigma(\sigma) \geq 0$ implies $\Sigma(\delta) \geq 0$. Since $|\delta| < n$, the induction hypothesis gives a decomposition $\xi_1 \ldots \xi_r$ for δ. But then $\eta\xi_1 \ldots \xi_r$ is the required decomposition of σ. □

Theorem 4.5. The set of reachable markings of a persistent Petri net is semi-linear.

Proof. Get $G = \{\xi_1, \ldots, \xi_n\}$ be the set of all firing sequences σ which satisfy $EPK(\sigma) \in F$. By Lemma 4.3, G is finite. For each $S = \{\xi_1^S, \ldots, \xi_{r_S}^S\} \subseteq G$, let A_S be the set of points $b \in N^{k+\ell}$ such that: 1) $b = (PK(\sigma), M)$ for some fireable firing sequence σ where $M_0 \xrightarrow{\sigma} M$ and 2) every sequence in S is fireable at M.

Let B_S be the finite set of minimal points of A_S. Let $\mathcal{M}_S$ be the finite set of markings which are associated with the points of B_S. We claim that the set of reachable markings is semi-linear and may be defined by

$$R = \bigcup_{S \subseteq G} \bigcup_{M \in \mathcal{M}_S} \{M + \sum_{i=1}^{r_S} c_i \Sigma(\xi_i^S) \mid c_i \geq 0, 1 \leq i \leq r_S\}.$$

First show that every marking in R is reachable. Fix $S \subseteq G$, $M \in \mathcal{M}_S$ and $\bar{M} = M + \sum_{i=1}^{r_S} c_i \Sigma(\xi_i^S)$. By definition, M is a reachable marking. Moreover, every member of S is fireable at M. But $EPK(\xi_i^S) \in F$ so $\Sigma(\xi_i^S) \geq 0$ for each $\xi_i^S \in S$ and hence each member of S is fireable an arbitrary number of times from M. Thus $\bar{M}$ is a reachable marking.

Let $\bar{M}$ be an arbitrary reachable marking and show that $\bar{M} \in R$. Let S be the set of sequences in G that are fireable at $\bar{M}$. Then $\bar{M}$ is part of a point $a = (PK(\gamma), \bar{M})$ in A_S. If $\bar{M} \in \mathcal{M}_S$, then we are done. If not, then there is an $M' < \bar{M}$ such that $M' \in \mathcal{M}_S$ is part of a minimal $b = (PK(\sigma), M')$, $b < a$, where $M_0 \xrightarrow{\gamma} \bar{M}$ and $M_0 \xrightarrow{\sigma} M'$ with $PK(\sigma) \times M' \in B_S$. Then $PK(\sigma) < PK(\gamma)$ and because the net is persistent, Lemma 3.1 implies that there is a fireable τ such that $PK(\tau) = \underline{max}(PK(\sigma), PK(\gamma)) = PK(\gamma)$ where $\tau = \sigma\alpha$ for some α. Hence $M_0 \xrightarrow{\tau} \bar{M}$ where $M_0 \xrightarrow{\sigma} M' \xrightarrow{\alpha} \bar{M}$. But $M' < \bar{M}$ implies that $\Sigma(\alpha) \geq 0$ so by Lemma 4.4 there is a sequence $\delta \in S^+$ such that $M' \xrightarrow{\delta} \bar{M}$. This implies that

$$\bar{M} = M' + \sum_{i=1}^{r_S} c_i \, \Sigma(\xi_i^S),$$

for some $c_i \geq 0$ where $S = \{\xi_1^1, \ldots, \xi_r^1\}$ so $\bar{M} \in R$. □

Theorem 4.5 shows that the set of reachable markings of a persistent net is semi-linear. We believe that this indicates why persistent nets have been more tractable than arbitrary Petri nets. In Section 6 we give examples which indicate that only a "small amount" of non-persistence allows us to simulate arbitrary Petri nets.

5. Conflict Free Nets

Recall that a Petri net is conflict free if each place p of the net either is an input for at most one transition or all such transitions are on self loops with p. In this section, we show that for conflict free nets the maximum number of tokens that any place can receive is either unbounded or is linearly bounded by the size of the initial marking. I.e., there is a contant c such that for an arbitrary initial marking with x tokens, the maximum marking of each place is either unbounded or is bounded by cx. The proof of this theorem yields an exponential time algorithm for deciding whether a conflict free net is bounded. The best known algorithm for deciding boundedness of arbitrary nets [3] requires a computing time which is bounded above by Ackermann's function.

Our first result relates persistent and conflict free nets.

Theorem 5.1. A Petri net is conflict free if and only if it is persistent for all initial markings.

Proof. $\Rightarrow$ obvious

$\Leftarrow$ Assume that a Petri net is not conflict free. Pick any two transitions t_1 and t_2 having an input place p in common where t_1 is not on a self loop with p. Put 1 token on each input place of t_1 and 1 token on each input place of t_2 which is not an input place of t_1. No other places are given tokens. Then both t_1 and t_2 are fireable but the firing of t_1 disables t_2 so the net is not persistent. □

Definition. A transition t of a of Petri net is ω-linearly bounded if there is a constant c such that for any initial marking having a total of x tokens, either: 1) t occurs an unbounded number of times in fireable firing sequences or 2) t occurs at most cx times in fireable firing sequences. A place p of a Petri net is ω-linearly bounded if there is a constant c such that for any initial marking having a total of x tokens, either p is unbounded or $M(p) \leq cx$ for any reachable marking M.

Lemma 5.2: Every transition of a conflict free Petri net is ω-linearly bounded.

Proof. The proof is by induction on the number of arcs from places to transitions. Let P be a conflict free Petri net. If P has no arcs from places to transitions, then the result is obvious, with $c = 1$ since each transition can fire an unbounded number of times.

Assume the theorem holds for all conflict free nets with n arcs as above and let P be a conflict free net with $n + 1$ arcs from places to transitions. Let t be an arbitrary transition of P. If t has no input places, the result is immediate. Assume that p is an input place of t in P. By the induction hypothesis, the theorem is true for the conflict free net P' obtained from P by deleting the arc from p to t. Let c be a constant which ω-linearly bounds the firings of the transitions of P' and let the initial marking of P have x tokens. We must show that the transition t of P is ω-linearly bounded. Because of the induction hypothesis, there are two cases:

Case 1. If t could fire at most cx times in P', then t can fire at most cx times in P.

Case 2. Assume t could fire an unbounded number of times in P'. If each of the c_1 transitions having p as an output place could fire at most cx times in P', then the same is true of these transitions in P. Then p receives at most $(c_1c+1)x$ tokens in P so t can fire at most $(c_1c+1)x$ times in P. Assume that some transition $\bar{t}$ having p as an output place could fire an unbounded number of times in P'. By Lemma 3.8, t and $\bar{t}$ occur infinitely often in some ω-firing sequence α of P'. By Lemma 3.10, α can be chosen so that

$$\alpha = \delta t \gamma_1 \bar{t} \delta_1 t \gamma_2 \bar{t} \delta_2 t \gamma_3 \cdots$$

where t does not occur in δ, γ_i and δ_i $(i \geq 1)$. Let $\alpha(k)$ be the initial segment of α up to and including the k-th occurrence of t. There are two subcases:

subcase a: If t cannot fire in P, then we are done.

subcase b: Let β be the shortest fireable sequence (possibly empty) of P which enables t. Since β is fireable in P, it is also fireable in P'. But then by Lemma 3.1, for any k, $\sigma_k = \beta\cdot(\alpha(k)-\beta)$ is fireable in P' and contains k occurrences of t. The i-th occurrence of t in σ_k is preceded by β plus at least $i - 1$ occurrences of $\bar{t}$ so before the i-th firing of t in σ_k (in P'), p contains at least i tokens. Thus implies that σ_k is fireable in P because P is obtained from P' by adding an arc from p to t. Hence in this case, t is unbounded in P. □

Definition. A loop in a Petri net is a non-trivial path, all of whose arrows point in the same direction, which begins and ends with the same place. A loop is unbounded if at least one place on the loop is unbounded. A transition is unbounded if it occurs an unbounded number of times on fireable firing sequences. A loop is immortal if every transition on the loop is unbounded.

Lemma 5.3. Consider a persistent net and a place p such that p is only on unbounded, immortal loops. Further assume that p is on such a loop. Then p is unbounded.

Proof. Assume that p is bounded. By Theorem 3.5, for each $i \geq 1$, there is a finite firing sequence which marks each unbounded place of the net with at least i tokens. Let k_i be the maximum p marking for such sequences. Then let k be the maximum value which occurs infinitely often among $\{k_i\}$ and let $\alpha(i_j)(j \geq 1, i_1 < i_2 < \ldots)$ be a finite firing sequence which marks each unbounded place with i_j tokens and p with k tokens. Choose a subset $\{\beta_1, \beta_2, \ldots\}$ of $\{\alpha(i_1), \alpha(i_2), \ldots\}$ such that $M_{i+1} > M_i$ for $i \geq 1$ where M_i is the marking after β_i has fired. Since p is on an immortal loop, by Lemma 3.1, there is an extension of β_1 which fires a transition t' which has p as an output place. Because $M_{i+1} > M_i$ for $i \geq 1$, the same extension can be appended to each β_i. But the definitic of k requires that (after β_i has fired) some t, having p as an input place must be fired before t' is fired. This means that p

is on a loop none of whose places, except p, contains a token after β_i has fired. (If some place, other than p, of every loop that contains p has a token, then t need not be fired to enable t'.) Because p is only on unbounded loops, this contradicts the definition of β_i. Therefore, p is unbounded. □

<u>Theorem 5.4</u>. Each place of a conflict free Petri net is ω-linearly bounded.

<u>Proof</u>. Let p be an arbitrary place of the net. Let c be the constant given by Lemma 5.2 which ω-linearly bounds the transitions of P. Assume that the initial marking of P has x tokens. There are four cases:

1. If all transitions into p are bounded, then the maximum marking on p is $(r+1)cx$ where there are r arrows into p.
2. If at least one transition into p is unbounded and is not on an immortal loop with p, then p is unbounded.
3. If all unbounded transitions into p are only on immortal, unbounded loops with p and there is at least one such loop, then p is unbounded by Lemma 5.3.
4. Assume that some unbounded transition into p is on a bounded, immortal loop with p. Then p is bounded by the definiton of boundedness for loops. The only way tokens can be added to the places of a loop in a conflict free net is by the firing of a transition not on the loop one of whose output places is on the loop. Moreover, the number of tokens on such a loop can never decrease. But then, because the loop is bounded, every transition into the loop may fire at most cx times so the maximum marking p can receive is $(r+1)x$ where r is the number of such transitions. □

A careful analysis of the proofs of Lemmas 5.2 and 5.3 and and Theorem 5.4 yields the following:

<u>Theorem 5.5</u>. There is an exponential time (in the size of the net) algorithm for deciding whether an arbitrary conflict free net is bounded.

Theorem 5.5 provides a substantial improvement over Karp and Miller's [3] algorithm for arbitrary nets which has Ackermann's function as an upper bound.

6. Examples and Conclusion

In this section we give some examples which illustrate the importance of persistence to the study of Petri nets. We also pose some open problems for future study.

Van Leeuwen's [9] reachability proof for 3-coordinate vector addition systems involves showing that such systems have semi-linear sets of reachable points. Since his methods do not generalize, the question arises as to how many coordinates (or places in a Petri net) are needed to generate non-semi-linear sets. Figure 3 gives a Petri net with 5 places whose reachability set is not semi-linear.

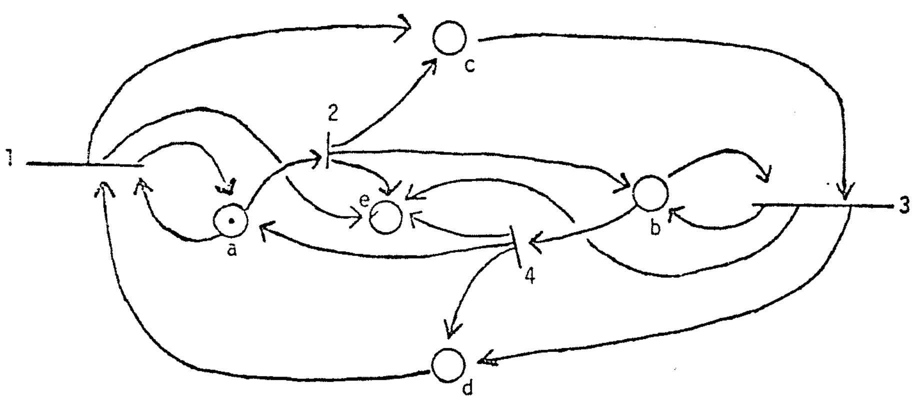

Figure 3

Any reachable marking M of the net of Figure 3 satisfies:

1. $M(a) + M(b) = 1$
2. $M(c) + M(d)$ = number of firings of 2 and 4
3. $M(e)$ = total number of transitions fired

We claim that if $M(c) + M(d) = n$, then the maximum value $M(e)$ can have is $\sum_{i=1}^{n+1} i - 1 = (n+1)(n+2)/2 - 1$. This is achieved as follows:


```
A:Fire 2
  Fire 3 until  c  has no tokens
  Fire 4
  Fire 1 until  d  has no tokens
  Goto A
```

Hence the set of reachable markings is not semi-linear. Notice that the net is conflict free (persistent) except for places a and b. Our next example will show that no additonal non-persistence is needed to achieve any set of reachable markings.

The net of Figure 3 translates to a vector addition system with seven coordinates. The two additional coordinates are needed because vector addition systems do not have self loops. An interesting problem, given Van Leeuwen's results, is to determine how many places (co-ordinates) are needed to generate non semi-linear sets.

The next three constructions illustrate the central role that persistence plays in Petri nets. We first show* that, in a very restricted sense, reachability is reducible to persistence. Together with the results of [10], this proves that persistence (in the restricted sense) is equivalent to reachability. It is important to note that this does not solve the difficult open problem regarding the equivalence of persistence and reachability. The paper concludes with a construction that shows that only a minimal amount of non-persistence is re uired to achieve any reachable set of markings.

Two transitions t_1 and t_2 are non-persistent if for some reachable marking M, $M \xrightarrow{t_1}$ and $M \xrightarrow{t_2}$ but not $M \xrightarrow{t_1 t_2}$.

Let P be an arbitrary Petri net. Let p be an arbitrary place of P. We show that if the non-persistence of two arbitrary transitions in a Petri net is decidable, then we can decide whether P has

*The statement of this result was communicated to us by M. Hack. The construction given is ours.

a reachable marking M such that $M(p) = 0$. To accomplish this, modify P as follows (Figure 4)

1. For each $t_i \in T$, add a new place p_i. Add p_i as an input place and as an output place of t_i. Initially p_i contains one token.

2. Add new transitions $t, \bar{t}, \bar{\bar{t}}$ and new places $\bar{p}, \bar{\bar{p}}$. Let $\{p_i | t_i \in T\}$ be the input set of t and let $\{p, \bar{p}, \bar{\bar{p}}\}$ be the output set of t. Let $\bar{t}(\bar{\bar{t}})$ have input places $\bar{p}$ and p (p and $\bar{\bar{p}}$).

Then $\bar{t}$ and $\bar{\bar{t}}$ are non-persistent iff $M(p) = 0$ for some reachable marking M.

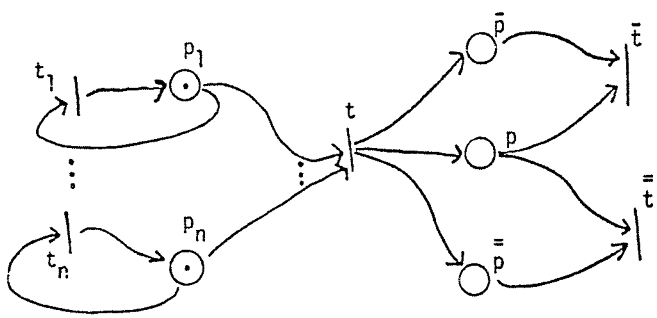

Figure 4

Reduction of reachability to non-persistence of two transitions

Let P be a Petri net which has r non-peristent transition where $2^{m-1} < r \le 2^m$, $r > 6$. Figure 5 shows how the number of non-persistent transitions can be reduced to 2m. By iterating this procedure, a new net is obtained with at most six non-persistent transitions (m non-persistent places).

The net P' obtained as in Figure 5 has additional places not occurring in P. However, if a marking M is reachable in P', then the restriction of M to the places of P is reachable in P. Conversely, if M is reachable is P, then some extension of M is reachable in P'. Similar results are true of fireable firing sequences.

Consider Figure 5 where $v_0, \ldots, v_{2^m}$ are the non-persistent transitions of the original net. Add transitions $t_1^0, t_1^1, t_2^0, t_2^1, \ldots t_m^0, t_m^1$ and places $b_1^0, b_1^1, b_2^0, b_2^1, \ldots, b_m^0, b_m^1, a_1, \ldots, a_m$ as shown in Figure 4 with each of $a_1, \ldots, a_m$ initially having one token. Now one of t_i^0 or t_i^1 $(1 \le i \le m)$ can fire, placing a token in b_i^0 or b_i^1 $(1 \le i \le m)$ respectively. Assume each v_j is indexed by an m place binary number.

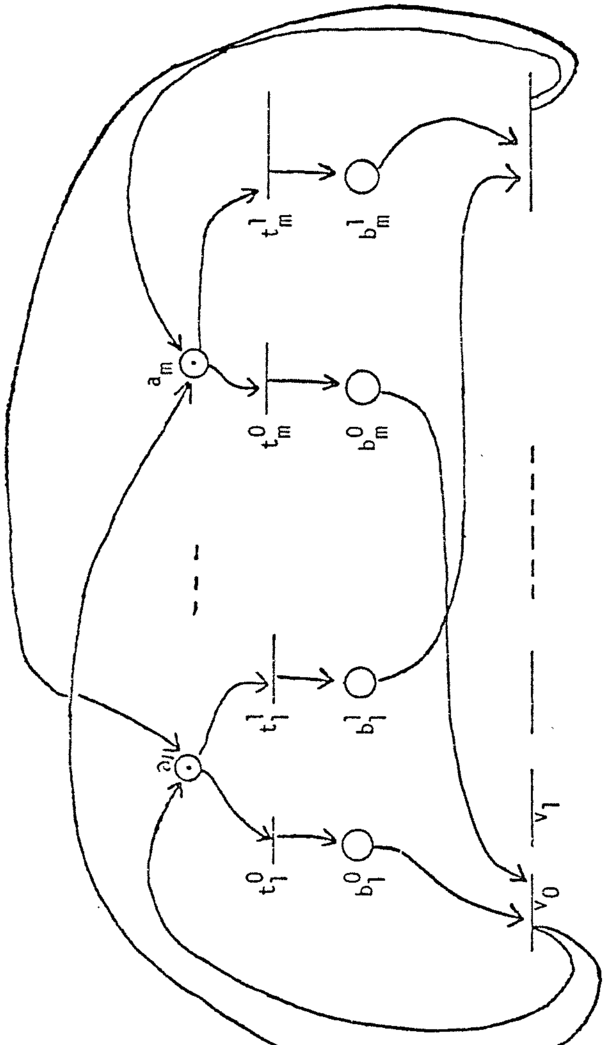

Figure 5

Reduction from 2^m to 2m non-persistent transitions

The firing of the combination of the t's whose superscripts correspond to this binary number will permit v_j to fire (when all inputs to v_j in the original net have tokens). After v_j fires, tokens are again placed in $a_1,\ldots,a_m$. Note that at most one of the v's can be enabled at any time and that that proper choice of t's will permit the enabling of any particular v_j. Hence the only non-persistent transitions in the modified net are the t's which have been added.

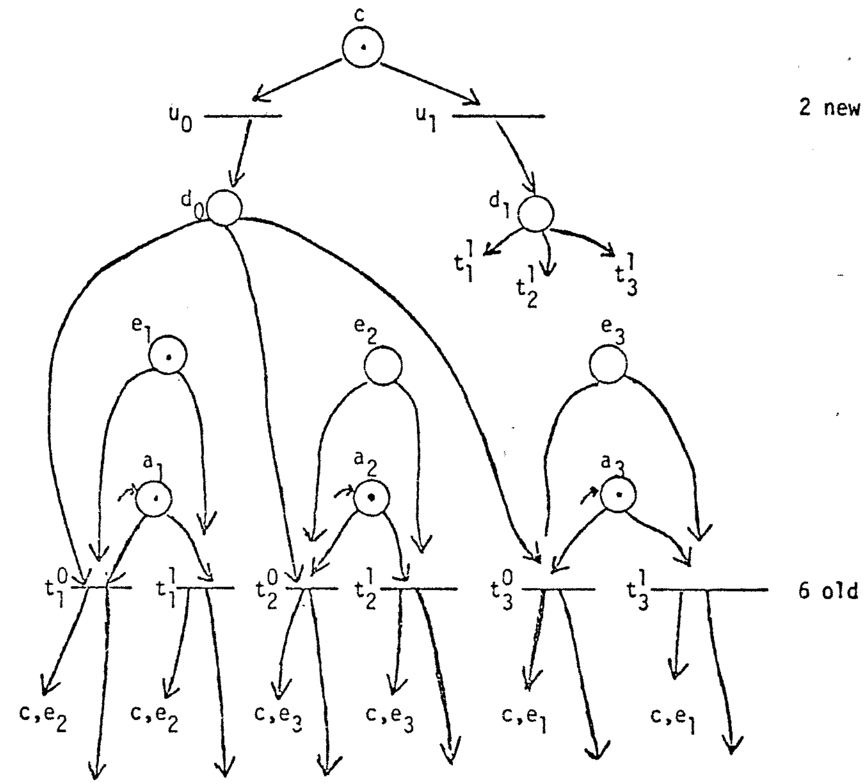

Figure 6

Reduce from 6 to 2 non-persistent transitions

Figure 6 illustrates how the six non-persistent transitions can be further reduced to two. The transitions $t_1^0, t_1^1, \ldots, t_3^1$ are the non-persistent transitions which result from iterating the process illustrated in Figure 5. The places a_1, a_2, and a_3 are as introduced in the last step of this process. One of t_1^0 or t_1^1 must fire before t_2^0 or t_2^1 can fire. Similarly t_2^0 or t_2^1 must fire before t_3^0 or t_3^1 can fire. Since only one of d_0 or d_1 can contain a token, only one of t_i^0 or t_i^1 $(1 \le i \le 3)$ can be enabled at any particular time. Therefore the t's can no longer lead to non-persistence. Only u_0 and u_1 are not persistent.

The results of this section show that all reachable sets (modulo a projection) can be realized by nets with only two non-persistent transistions. This is particularly interesting in light of the results of Section 4 that persistent nets only generate semi-linear sets. We believe that a thorough understanding of persistence will be necessary if the reachability problem is to be solved. Some interesting (and difficult) open problems are:

1. Is reachability decidable for persistent nets?
2. Is there an algorithm which decides whether an arbitrary net is persistent?
3. Is reachability decidable for arbitrary nets?
4. It has been shown that persistence is reducible to reachability [10]. Is the converse true?
5. Can Lipton's exponential time lower bound for the reachability problem be improved?
6. Can the Karp-Miller algorithm for deciding boundedness be improved?

Acknowledgment: We gratefully acknowledge several stimulating conversations with P. S. Thiagarajan and in particular his suggestion that semi-modular lattices might be important in studying persistent Petri nets.

References

1. Crespi-Reghizzi, S. and Mandrioli, D. A decidability theorem for a class of vector-addition systems, Inf. Proc. Letters 3(1975), 78-80.

2. Hack, M. The recursive equivalence of the reachability problem and the liveness problem for Petri nets and vector addition systems, 15 Annual IEEE SWAT Symposium, 1974, 156-164.

3. Karp, R. M. and Miller, R. E. Parallel program schemata, JCSS 3(1969), 147-195.

4. Keller, R. M. Vector replacement systems: a formalism for modelling asynchronous systems, TR 117, Dec. 1972, Princeton Univ. Computer Science Lab.

5. Lipton, R. J. The reachability problem and the boundedness problem for Petri nets are exponential-space hard, Conf. on Petri Nets & Related Methods, MIT, July, 1975.

6. Lipton, R. J., Miller, R. E. and Snyder, L. Synchronization and computing capabilities of linear asynchronous structures, 16th Ann. IEEE Symp. on Foundations of Computer Science, 1975, 19-28.

7. Parikh, R. J. On context free languages, JACM 13(1966), 570-581.

8. Petri, C. A. Fundamentals of a theory of asychronous information flow, Proc. IFIP Conf., 1962, Munich, 166-168.

9. Van Leeuwen, J. A partial solution to the reachability problem for vector addition systems, 6th Ann. ACM Symp. on Theory of Computing, May 1974, 303-309.

10. Jones, N., Landweber, L. H., and Lien, Y. E. Complexity of Some Petri Net Problems, in preparation.