Algorithms on marked directed graphs
Algorithmes sur graphes orientés définis

By M.A. Comeau *and* K. Thulasiraman, *Faculty of Engineering, Concordia University, Montreal.*

Marked directed graphs are a special case of Petri nets/ introduced by Carl Petri as a model for information flow in systems exhibiting asynchronism and parallelism. Commoner, Holt, Even and Pnueli* have studied several structural and algorithmic aspects of marked graphs using graph theory and network flow algorithms. Subsequently, Murata* has studied these graphs using a circuit-theoretic approach. In this paper, we combine the ideas of both these works and present algorithms for certain problems on marked graphs. In particular, we introduce the concept of scatter in firing sequences and present algorithms for determining minimum scatter firing sequences for different classes of graphs. We also present an approach for the general case and conclude with a lowerbound on the scatter of any firing sequence leading from an initial marking to a reachable final marking on a given marked graph.

Les graphes orientés définis constituent un cas spécial des Petri nets/ introduits par Carl Adam Petri comme un modèle de circulation de l'information dans les systèmes présentant des caractéristiques d'asynchronisme et de parallélisme. Commoner, Holt, Even et Pnueli* ont étudié plusieurs aspects structurels et algorithmiques de graphes définis en employant la théorie des graphes et des algorithmes de débit de réseau. Par la suite, Murata* a étudié ces graphes en utilisant une approche théorique sur les circuits. Dans cette étude, nous conjuguons les idées de ces deux travaux et présentons des algorithmes pour certains problèmes sur des graphes définis. En particulier, nous présentons le concept de la dispersion dans les séquences d'amorçage et présentons des algorithmes pour déterminer les séquences d'amorçage à dispersion minimum pour différentes classes de graphes. Nous présentons également une approche pour le cas général et concluóns par une limite inférieure sur la dispersion de toute séquence d'amorçage menant d'une marque initiale à une marque finale réalisable sur un graphe défini donné.

Introduction

A Petri net[1] is a general abstract algebraic structure originally developed by Carl Adam Petri as a model for information flow in systems exhibiting asynchronism and parallelism. Petri net modeling has applications in computer communications, operating systems, operations research, artificial intelligence, as well as in physiological models of the brain. The generality of the Petri net makes modeling of complex networks possible. However, the feasibility of analysis becomes questionable and, in many cases, the problems are NP-complete, with solutions sometimes undecidable. The Petri net is a bipartite graph structure in which there are two groups of nodes called transitions and places, and an edge set connecting transitions to places and vice versa.

Marked directed graphs are a restricted class of Petri nets in which only transition nodes are present. The places are absorbed in the edge set of the graph. Thus, a Petri net reduces to a marked directed graph if every place in the Petri net has exactly one input transition and one output transition. Being a special class of Petri nets, marked directed graphs are more amenable to analysis, yet they retain enough generality to model systems of parallel processing, queueing networks, resource allocation schemes and many other related problems. Existing analysis techniques are essentially based on graph theoretic properties of the underlying directed graphs. This paper combines a circuit theoretic approach following that of Murata[2] with an algorithmic approach following Commoner, Holt, Even and Pnueli[3] to present algorithms for certain problems on marked directed graphs. In particular, the following problems are considered: executability of a firing count vector and generation of a minimum scatter firing sequence for a given executable firing count vector. First, the basic definitions and terminology are established.

A marked directed graph is a directed graph G with a marking M defined as a nonnegative integer vector associated with the edge set of the directed graph. We assume that G is finite. If G has e edges, then M has e components—one for each edge of G. The components of M are referred to as token counts for the corresponding edges. An action or event takes place in a marked directed graph by firing a vertex or node. Firing a node changes the token count of all edges incident on the node. All input edges lose one token and all output edges gain one token. Since the marking must remain nonnegative after any node firing, it is never possible to fire any node with a token-free input edge. Thus, to be legally fired, a node has at least one token on each of its input edges. The firing of such a node is called a legal firing. If the firing of a node represents the activation of some process and the input edges represent necessary conditions for the process to occur, then a node which can be legally fired is equivalent to a process with all of its input conditions satisfied. Here, tokens represent flags or Boolean variables indicating whether or not a condition is satisfied. Thus, multiple tokens on an edge in a marked directed graph may be redundant or meaningless in some models. In such cases, node firing can be redefined, or topological constraints can be imposed on G, so as to restrict the token count on every edge to a maximum of one. Marked directed graphs with this property are called safe.

A marking M_n is said to be reachable from a marking M_0 if a legal firing sequence exists which transforms M_0 to M_n. Consider any two markings M_0 and M_n of a marked directed graph G such that M_n is reachable from M_0. Let $\Delta M = M_n - M_0$. We shall refer to ΔM as the differential marking. Murata has shown[2] that ΔM satisfies KVL equations. Thus, if B_f is the fundamental circuit matrix of G, then $B_f \Delta M = 0$. So, we can consider the elements of ΔM as the voltages of the corresponding edges of G. Using well-known, network-theoretic results, we can determine the node voltages $\sigma_1, \sigma_2, \ldots$ of G such that $\sigma_i - \sigma_j = \Delta M(e)$ where e is the edge directed from vertex i to vertex j and $\Delta M(e)$ is the component of ΔM corresponding to e.

Let Σ denote the column vector of σ_i's. We shall refer to Σ as the firing count vector. If Σ has any negative elements, the same number can be added to all the elements of Σ so that a smallest entry is a zero. The firing count vector Σ_0 obtained by this operation is called the minimum firing count vector. Clearly Σ_0 is unique, if it exists. The i^{th} element of an executable Σ_0 denotes the number of times vertex i would fire in a firing sequence leading from M_0 to M_n. The vertex with zero firing count will be referred to as a datum. The existence of a firing count vector Σ, satisfying KVL, does not guarantee the existence of a legal firing sequence leading from M_0 to M_n. A firing count vector Σ is said to be executable from M_0 if a legal firing sequence exists starting from M_0 and its firing count vector is Σ.

Executability of a firing count vector

In this section we give an algorithmic proof of the following theorem due to Murata.[2] The proof here is based on that of Theorem 5.[3]

Theorem 1

Let M_0 and M_n be any two markings of a marked directed graph G. Let $\Delta M = M_n - M_0$. A minimum firing count vector $\Sigma_0 = [\sigma_i]_{v \times 1}$ satisfying

$$A^t \Sigma_0 = \Delta M, \quad (1)$$

where v is the number of nodes of G and A is the incidence matrix of G, is executable from M_0 if and only if $\sigma_k = 0$ for each vertex k of every token-free directed circuit at M_0.

Proof

Necessity: obvious.

Sufficiency: we may assume that $\Sigma_0 \neq 0$. Otherwise, $M_0 = M_n$ and the case is trivial. To prove the theorem, we shall first show that there exists a legally firable vertex with its firing count nonzero. Consider any $\sigma_{i_1} > 0$. If vertex i_1 is not legally firable, then there exists an edge $e = (i_2, i_1)$ directed from some vertex i_2 to vertex i_1 with $M_0(e) = 0$. If vertex i_2 is also not firable, then repeat this process until a sequence of vertices $i_1, i_2, \ldots, i_{k-1}, i_k$ is located such that:

- $i_1, i_2, \ldots, i_{k-1}$ are all distinct;
- $\sigma_{i_k} > 0$ for $j = 1, 2, \ldots, k$;
- for each edge $e = (i_{j+1}, i_j)$, $j = 1, 2, \ldots, k-1$, directed from vertex i_{j+1} to vertex i_j, $M_0(e) = 0$;
- Case 1: vertex i_k is legally firable or
  Case 2: $i_k = i_j$, for some $j = 1, 2, \ldots, k-2$.

One of the two Cases above should occur because the graph G is assumed to be finite. If Case 1 occurs, then a legally firable vertex with nonzero firing count has been found. Case 2 cannot occur for that would mean the existence of a token-free directed circuit i_{k-1}, $i_{k-2}, \ldots, i_j, i_{k-1}$ with each of the vertices on this circuit having a nonzero firing count, thereby contradicting the hypothesis of the theorem.

Now, fire the vertex i_k. This would result in a new marking M_1 and a new firing count vector Σ_1. The marking M_1 is obtained by increasing by one the token count of all edges directed away from i_k, and by decreasing by one the token count of all edges directed into i_k. Also, Σ_1 is obtained from Σ_0 by subtracting one from σ_{i_k}. Clearly, M_1, Σ_1, and the vector $M_n - M_1$ satisfy the hypothesis of the theorem. If $M_1 = M_n$, then $\Sigma_1 = 0$ and we have established the executability of Σ_0 starting at M_0. Otherwise, locate a legally firable vertex with respect to M_1 and fire this vertex. Repeat this process until a marking is reached equal to M_n and with a firing count vector Σ_n equal to zero. The hypothesis of the theorem guarantees the termination of this process and hence the executability of Σ_0 starting at M_0 leading to M_n. //

Scatter in a firing sequence

A firing traversal of a marked directed graph is a traversal of the nodes of the graph which executes a given legal firing sequence. Thus, a firing traversal must visit nodes of a marked directed graph in the order dictated by the legal firing sequence, and update the marking by firing the nodes. With this algorithmic definition of a firing traversal, equivalent firing sequences leading from a common initial marking to the same final marking may be studied. Consider executing the firing sequence $a^2b^3ac^4$. Using this representation, the node labeled a is visited and fired twice, then node b is visited and fired three times, etc. Among all possible legal firing sequences between two markings on a marked directed graph, which of these will require a minimum or maximum number of node visits in a firing traversal of the marked directed graph? This leads to the concept of scatter in a firing sequence.

The scatter of a firing sequence F, denoted by scatter (F), is the difference between the number of node visits needed to execute F, and the number of distinct nodes visited in a firing traversal. For example, if $F = a^2b^3ac^4$, then scatter $(F) = 1$ because a firing traversal with F would require four node visits (a, b, a, c) and three distinct nodes (a, b, c) would be visited. Thus, the scatter is a non-negative integer.

The definition of scatter applies to a compressed representation of firing sequences. That is, consecutive elementary firings of a node are considered to occur during the same visit of that node. An algorithm to perform a firing traversal might use a representation of F as an ordered list of pairs of the form (label, coefficient). The scatter would then be the number of pairs in a firing sequence minus the number of distinct labels. If the scatter of a firing sequence is zero, then a firing traversal can execute it by visiting each node to be visited exactly once. For a general graph, a zero scatter firing sequence is not always present between an intitial marking and a reachable final marking. This leads to the problem of finding a minimum scatter, legal firing sequence leading from some initial marking to a reachable final marking.

We now discuss an application which provides the motivation for introducing the concept of scatter. Consider modeling a general industrial environment with a marked graph as shown in Figure 1. Each node in this graph represents an activity. In this figure, nodes 1 to k represent sources of possibly different commodities being introduced into the environment from the external world. Nodes 1′ to p' represent commodity sinks or distribution centers to the external world. The token count on edge (i,j) incident into node j

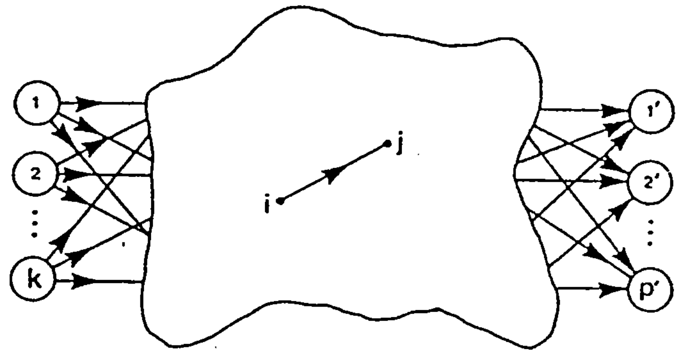

Figure 1: *Modeling with a marked graph.*

represents the number of units of some commodity made available for processing at node j. For an activity i to occur, there must be at least one unit of all the required commodities, which in fact corresponds to having at least one token on each input edge of node i. Occurrence of an activity i thus corresponds to firing node i in the marked graph. Firing a node consumes one unit of commodity from each input edge and produces one unit of commodity on each output edge. If edge (i,j) is incident out of node i, then completion of activity i makes one unit of some commodity available for processing at node j. It may be noted that initially certain edges may have a non-zero token count. This may represent the initial state of the environment. The final state corresponds to a pre-defined level of production, made available for consumption to the external world through the distribution centers. In other words, in the final state, the edges incident into the sinks will be required to have a certain specified number of tokens.

At each node, a "machine" is available for carrying out the corresponding activity. Starting a machine from the shut down state may involve considerable overhead cost. It is obviously cost-effective to minimize the idle running times of the machines, so a machine at any node is shut down whenever that node cannot be activated. Thus, it is desirable to minimize the number of machine starts. It should be easy to see that the length of the firing sequence leading from the initial marking to the final marking is a measure of the overhead cost incurred in reaching the desired production level. This motivates the introduction of the concept of scatter as well as the problem of determining minimum scatter firing sequences. This model can be further generalized by modifying the firing rule or by incorporating a set of firing rules or semantics[4] as described for the design of a distributed operating system kernel.

To illustrate the concept of scatter in firing sequences, consider the markings M_0 and M_n of the directed graph shown in Figure 2. The firing count vector $\Sigma_0 = [\sigma_a, \sigma_b, \sigma_c, \sigma_d]^T = [3, 3, 4, 0]^T$ is ex-

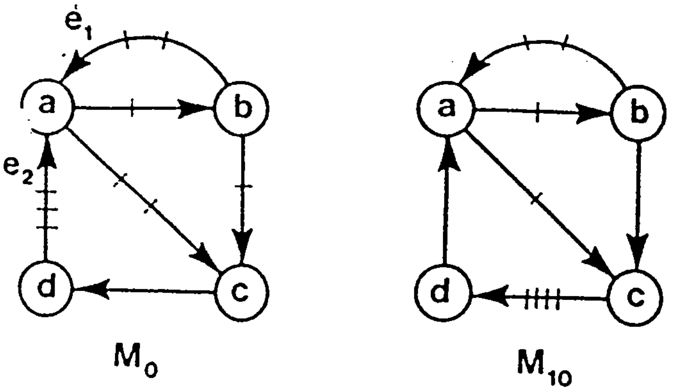

Figure 2: Example illustrating scatter of a firing sequence.

ecutable from M_0 and leads to M_*. Some possible legal firing sequences are $a^2b^3ac^4$, $ba^3b^3c^4$, $ca^2b^3ac^3$, and $ab^2a^2bc^4$. The first two of these scatter once and the last two scatter twice. The sequences $a^2b^3ac^4$ and $ba^3b^3c^4$ are the minimum scatter firing sequences executing Σ_0 from M_0. All other legal firing sequences executing Σ_0 from M_0 scatter at least twice. This can be easily verified for this simple graph by obtaining all possible legal sequences from M_0 to M_* which execute Σ_0.

The enabling number of a node in a marked directed graph is the number of times it can be legally fired when visited at a given visit in a firing traversal if this number is less than or equal to the node's firing number. Otherwise, the enabling number of a node is its firing number. The disabling number of a node in a marked directed graph is the difference between the node's firing number and its enabling number.

Before presenting the algorithms for determining minimum scatter firing sequences, a few observations can be made about marked ...cted graphs.

- Since a datum node always exists, at least one firing count is zero. That is, at least one node of the graph need never be fired to bring any initial marking to a reachable final marking.
- The markings on all input edges to the datum node must either remain the same or strictly increase from their initial values in any legal firing sequence.
- The markings on all output edges from the datum node must either remain the same or strictly decrease from their intitial values in any legal firing sequence.
- If some node is encountered that is enabled to its corresponding firing number at any visit during the generation of a firing sequence, it can be fired repeatedly until its firing number has been satisfied and then removed from the graph together with all edges incident on it. The problem then reduces to finding the remainder of the firing sequence on the resulting subgraph.
- Any source node, (that is a node with no input edges) can be fired independently since it has no input edges. Thus, a source need only be visited once in a firing traversal.

The greedy firing policy

In the example of the previous section, the enabling number of node a under the marking M_0 is

$$\mu_a = \min\{\sigma_a, M_0(e_1), M_0(e_2)\} = 2. \tag{2}$$

There are two possible ways of firing node a if visited under the marking M_0. The firing sequences $ab^2a^2bc^4$ and $a^2b^3ac^4$ are examples of both. In general, if μ_i is the enabling number of node v_i under a marking M, then ther are μ_i possible ways to fire node v_i if ...ed under M (since it can be fired up to μ_i times and it must be fired at least once to qualify as a visit).

An algorithm that generates minimum scatter firing sequences must employ some node firing policy at each node visit in order to alleviate the ambiguity associated with multiply-enabled nodes. An obvious choice is to fire each node as much as possible at each visit. Before employing such a greedy firing policy, we must show that we do not overlook all optimal solutions by examining only those sequences which fire a node to its enabling number at each visit. In other words, we must prove that a minimum scatter, legal firing sequence exists between any reachable markings, on a marked directed graph, in which each visit fires some node to its enabling number at that visit of the node. Such a firing sequence will be referred to as a greedy firing sequence. Unless explicitly stated, all firing sequences referred to will be assumed legal.

Theorem 2

For any executable firing count vector Σ_0, from a marking M_0 on a marked directed graph G, there exists a greedy, minimum scatter firing sequence F_g, executing Σ_0 from M_0, in which each visit of a node fires the node to its corresponding enabling number.

Proof

Let F be any minimum scatter sequence executing Σ_0 from M_0, and let m denote the total number of visits in F. If F is not a greedy sequence then let v_i be the first node of G which F fails to fire to its enabling number at the r^{th} visit. Construct another legal firing sequence F_g^r which is identical to F in the first $r-1$ visits and which fires node v_i to its enabling number μ_i^r at the r^{th} visit. The sequence F_g^r is r visits long and greedy at each visit. The enabling numbers of all nodes of G except node v_i after F_g^r executes from M_0, are greater than or equal to those after the first r visits of F from M_0. Thus, if F fires node $v_j, j \neq i$, at the $(r+1)^{th}$ visit then we may legally construct the firing sequence F_g^{r+1} defined by the relation

$$F_g^{r+1} = F_g^r v_j^{\mu_j^{r+1}}, \tag{3}$$

where μ_j^{r+1} is the enabling number of node v_j at the $(r+1)^{th}$ visit. The firing sequence F_g^{r+1} visits nodes in the same order as the first $r+1$ visits of F and is greedy at each visit. Since the node v_i is not fired to its enabling number at the r^{th} visit in F, it follows that v_i will be visited at least once after the r^{th} visit in F. Clearly, $v_j \neq v_i$. So, F must be at least $r+2$ visits long. In other words, $m \geq r+2$.

Now, two cases may arise; F either returns to node v_i or visits some other node v_k, $k \neq i, j$ at the $(r+2)^{th}$ visit.

Case 1: If F returns to v_i at the $(r+2)^{th}$ visit, then firing node v_j at the $(r+1)^{th}$ visit must have increased the enabling number of node v_i, since by hypothesis, F has minimum scatter and, therefore, does not visit nodes redundantly. So, the greedy firing sequence F_g^{r+1} which left node v_i disabled after the r^{th} visit must have re-enabled it after firing node v_j at the $(r+1)^{th}$ visit and, therefore, it is legal to construct the sequence F_g^{r+2} defined by the relation

$$F_g^{r+2} = F_g^{r+1} v_i^{\mu_i^{r+2}}, \tag{4}$$

where μ_i^{r+2} is the enabling number of node v_i at the $(r+2)^{th}$ visit.

Case 2: If F fires node $v_k \neq v_j, v_i$ at the $(r+2)^{th}$ visit, then by the previous argument, we may legally construct the sequence F_g^{r+2} defined by the relation

$$F_g^{r+2} = F_g^{r+1} v_k^{\mu_k^{r+2}}, \tag{5}$$

where μ_k^{r+2} is the enabling number of node v_k at the $(r+2)^{th}$ visit.

In either case, F_g^{r+2} visits nodes in the same order as the first $r+2$ visits of F and is greedy at each visit. By repeated application of the above construction, we will arrive at a sequence $F_g^m = F_g$ executing Σ_0 from M_0 in which each visit of a node fires the node to its corresponding enabling number.//

Since the construction described in the proof of the above theorem applies to any minimum scatter firing sequence executing Σ_0 from M_0, we have the following corollary.

Corollary 2.1

For every mimimum scatter, legal firing sequence F, executing Σ_0 from M_0 on a marked directed graph G, there is a greedy, minimum scatter, legal firing sequence F_g executing Σ_0 from M_0 which visits nodes in the same order as F does.//

This theorem implies that if we employ a greedy firing policy at each node visited when executing some firing count vector, then it is possible to execute that vector with minimum scatter by visiting the nodes properly. In other words, we cannot obtain a firing sequence with less scatter by leaving a node enabled after visiting and firing it, than we can by firing it to its enabling number at that visit. However, being less than greedy at each visit when firing nodes may lead to unnecessary scatter while searching for firing sequences, as illustrated by the firing sequences $ab^3a^3bc^4$ and $a^3b^3ac^4$, for the marked directed graph of Figure 2. Therefore, this is sufficient grounds for employing the greedy firing policy at each node while searching for a minimum scatter firing sequence.

The firing sequence $ca^2b^3ac^3$ for the marked directed graph of Figure 2 shows that arbitrarily applying the greedy node firing policy over the nodes of a graph may not produce minimum scatter firing sequences. Thus, along with the greedy node firing policy, some node visiting policy is needed to characterize an algorithm which generates a minimum scatter legal firing sequence. We pursue this further for general graphs. First, we obtain results for marked graphs with different topological properties.

Minimum scatter legal firing sequences

This section examines the problem of determining a minimum scatter firing sequence executing a given executable firing count vector from a given initial marking on marked graphs with different topological complexities. It is shown that any executable firing count vector, from a given initial marking on a marked acyclic directed graph, always possesses a zero scatter firing sequence. This is shown to be true for the simple directed circuit but not for a general topology with directed circuits. Disjoint directed circuits are then examined and an algorithm is given for each case.

Marked acyclic directed graphs

By definition, an acyclic directed graph has no directed circuits. Thus, on a marked acyclic directed graph G, a marking M_n is reachable from an initial marking M_0 if and only if KVL is satisfied between the markings. In an acyclic directed graph, there is at least one node with zero in-degree and at least one node with zero out-degree. A marked directed graph with the acyclic property must, therefore, possess at least one source and at least one sink. A source is independently firable since, by definition, it has no input edges. Thus, a firing traversal need only visit sources once.

The acyclic property of a marked acyclic directed graph can be used to generate a minimum scatter firing sequence between two markings with an executable firing count vector. If a source is located, it can be fired to its firing number and removed from the graph. This source is guaranteed in acyclic graphs. After removing this source from the graph, the subgraph thus induced must also have the acyclic property and the procedure can then be recursively applied until all nodes have been removed. This is just a topological sort[5] applied to the nodes of the acyclic graph. Let $v_1, v_2, \ldots, v_v$ be the labels assigned to the nodes of a marked acyclic directed graph G according to a topological sort, then

$$v_1^{\sigma_1} v_2^{\sigma_2} \ldots v_v^{\sigma_v} \tag{6}$$

is a zero scatter firing sequence for a given executable firing count vector Σ. So, we have the following theorem.

Theorem 3

A zero scatter, legal firing sequence exists for any given executable firing count vector on any marked acyclic directed graph.//

Figure 3 illustrates the procedure for a marked acyclic directed graph. The initial marking is irrelevant during topological sort, so a minimum scatter sequence can be obtained from the underlying graph and the firing count vector.

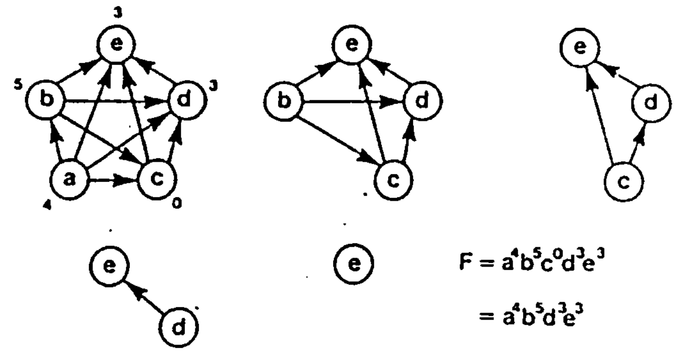

Figure 3: *A marked acyclic graph and a minimum scatter sequence determination.*

The marked directed circuit

We next consider the problem of generating a minimum scatter firing sequence between two markings on a marked directed circuit with a legally executable, minimum firing count vector. Let v be the number of vertices of a simple marked directed circuit C, pick any vertex of C and label it v_0. Traverse the circuit from vertex v_0 in the circuit direction labeling the vertices $v_1, v_2 \ldots, v_{v-1}$. Also, since the circuit has v edges, let e_j denote the edge incident into vertex v_j, for $j \in \{0, 1, \ldots, v-1\}$. Let M_0 and M_n be the initial and final markings of C and let Σ_0 be the minimum firing count vector. Since Σ_0 is assumed executable, KVL is satisfied between M_0 and M_n and thus

$$\sum_{j=0}^{v-1} M_0(e_j) = \sum_{j=0}^{v-1} M_n(e_j) = \Gamma_C \tag{7}$$

where Γ_C is the circuit token count. So, KVL for the marked directed circuit implies that any two markings on C are mutually reachable if and only if they have the same circuit token count.

Theorem 4

A zero scatter firing sequence exists between any two reachable markings on a simple marked directed circuit.

Proof

Observation 1 in the third section guarantees the existence of at least one vertex with a zero firing count. Observations 2 and 3 indicate that any datum vertices can be removed from the graph. If this datum vertex is removed from the circuit C, the subgraph induced is acyclic and Theorem 3 guarantees the existence of a zero scatter firing sequence by applying topological sort on it.//

Specifically, the firing sequence defined by

$$F = \prod_{j=0}^{v-1} v_{(i+j)\bmod(v)}^{\sigma_{(i+j)\bmod(v)}}, \tag{8}$$

where i is the index of any datum vertex, is always executable and has zero scatter. This sequence is obtained by starting at a datum vertex and sorting the circuit vertices in the order dictated by the circuit. Any nodes with zero firing count can be neglected. Figure 4 illustrates the construction of such a sequence.

Vertex-disjoint marked directed circuits

To extend the result for a simple marked directed circuit to the case of disjoint circuits, we first note that the existence of a datum node in the circuit is not guaranteed if other nodes are present in the graph. Thus, it may be necessary to fire every node of a directed circuit in a marked directed graph. Since the nodes of a directed circuit C are restricted to at most Γ_C firings per visit, scatter may occur. As an immediate consequence of Theorem 3, directed circuits are the only cause of scatter in a minimum scatter firing sequence.

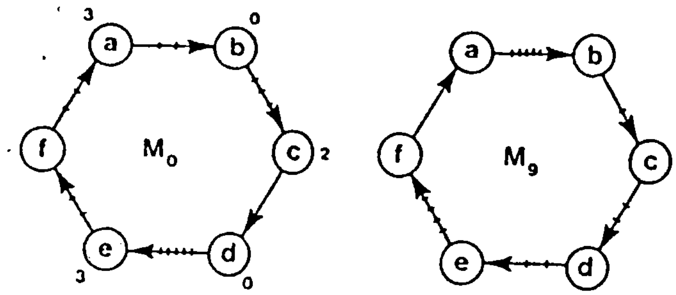

$$F_1 = b^0c^2d^0e^3f^1a^3 = c^2e^3f\,a^3$$

$$F_2 = d^0e^3f^1a^3b^0c^2 = e^3f\,a^3c^2$$

Figure 4: *A simple directed circuit and a minimum scatter sequence determination.*

Assuming that the vertices of a marked directed circuit C, which is vertex-disjoint with any other directed circuits in a marked directed graph G, are fire-restricted only by the markings on the circuit edges, we can analyze C independently for a minimum scatter subsequence. All vertices of G, except those of C, can be removed. The resulting subgraph is C with a corresponding set of firing numbers. One of two cases must occur. Either a vertex of C exists with a zero firing number or it does not. The former case reduces to the previous problem and zero scatter is possible in the subsequence for C by removing the vertex with zero firing number and topologically sorting the induced acyclic graph. In the latter case, a zero scatter subsequence may still be possible but this is not generally true.

At this point, the question of scatter in a subsequence for the circuit can be answered by examining the disabling numbers of the circuit vertices. If any disabling number is zero, then a zero scatter subsequence is possible starting at that vertex. This follows from the fact that after this vertex is fired to its enabling number, its firing number becomes zero and it can be removed from C, resulting in the acyclic case discussed previously. We are left with the problem of finding a minimum scatter firing sequence of C when the scatter is nonzero.

Where used as a vertex or edge index, $(i+j)$ will denote addition of the integers i and j, modulo v, where v is the number of vertices of C. Also, (j) will represent j mod v. With C labelled as in the previous subsection, we define a cyclic firing sequence of C, of length m, as one with the form

$$\prod_{j=1}^{m} v_{(k+j-1)}^{a_j} \tag{9}$$

where $a_j > 0$, for $j = 1, 2, \ldots, m$. Expressing m as $sv + r$, where $s = \lfloor m/v \rfloor$, and $r = (m)$, the cyclic sequence can be written as

$$v_k^{a_1} v_{(k+1)}^{a_2} \ldots v_{(k+v-1)}^{a_v} v_k^{a_{v+1}} \ldots v_{(k+r-1)}^{a_{sv}} \ldots v_{(k+r-1)}^{a_{sv+r}}. \tag{10}$$

and has scatter $m - v$ if $m \geq v$.

Theorem 5

Every minimum, nonzero scatter firing sequence F_C of a disjoint directed circuit C in a marked directed graph G, is cyclic.

Proof

Consider an arbitrary minimum, nonzero scatter firing sequence F_C of the vertices $v_0, v_1, \ldots, v_{v-1}$ of a disjoint directed circuit C in a marked directed graph G. The vertices are assumed to be cyclically labeled. Let F_C be written as

$$F_C = \prod_{j=1}^{m} x_j^{a_j} \tag{11}$$

where x_j is the vertex visited at the j^{th} visit and a_j is the number of times x_j is fired at the j^{th} visit. Since F_C has nonzero scatter, let v_k be the first vertex of C which repeats in F_C. let p and t denote the first and second visits of v_k in F_C, respectively. The following statements hold for F_C.

- $x_p, x_{p+1}, \ldots, x_{t-1}$ are distinct.
- $t - p \leq v$.
- $v_{(k-1)} \in \{x_j\}_{j=p+1}^{t-1}$.

The first statement follows from the hypothesis that v_k is the first vertex which repeats in F_C. The second statement follows from statement 1 and the fact that there are only v vertices in C. The last statement must be true if F_C is a minimum scatter firing sequence; otherwise, it is possible to construct another legal firing sequence F_C' by absorbing the t^{th} visit into the p^{th} visit, resulting in a sequence which has less scatter than F_C.

Identify the maximal set of visits $\{t_r\}_{r=1}^{s}$ satisfying the following properties:

a: $p + 1 \leq t_r \leq t - 1$, $r \in \{1, 2, \ldots, s\}$,

b: $x_{t_r} = v_{(k-r)}$, $r \in \{1, 2, \ldots, s\}$,

c: $t_{r+1} < t_r$, $r \in \{1, 2, \ldots, s-1\}$.

Thus, F_C has the form

$$F_C = x_1^{a_1} \ldots x_{p-1}^{a_{p-1}} v_k^{a_p} \ldots v_{(k-s)}^{a_{t_s}} \ldots v_{(k-2)}^{a_{t_2}} \ldots v_{(k-1)}^{a_{t_1}} \ldots v_k^{a_t} x_{t+1}^{a_{t+1}} \ldots x_m^{a_m}. \tag{12}$$

Two cases of interest are considered.

Case 1: $(k-s) \neq (k+1)$.
In this case, a_t is independent of the firing of v_k at the p^{th} visit. Thus, we may construct another legal firing sequence F_C' by concatenating the visits $\{t_r\}_{r=1}^{s}$ in the order in which they appear in F_C and placing them immediately before the p^{th} visit. After this rearrangement, the p^{th} and t^{th} visits may be grouped together as $v_k^{a_p+a_t}$ at the $(s+p)^{th}$ visit of F_C'. The constructed sequence F_C' will have the form

$$F_C' = x_1^{a_1} \ldots x_{p-1}^{a_{p-1}} v_{(k-s)}^{a_{t_s}} \ldots v_{(k-2)}^{a_{t_2}} v_{(k-1)}^{a_{t_1}} v_k^{a_p+a_t} \ldots x_{t+1}^{a_{t+1}} \ldots x_m^{a_m}. \tag{13}$$

and is $m - 1$ visits long. This construction would reduce the scatter in F_C but since F_C has minimum scatter, this is not possible. In other words, case 1 cannot occur.

Case 2: $(k-s) = (k+1)$; i.e. $s = v - 1$.
In this case, statement 2 holds with equality. That is, $t - p = v$ and so $\{x_{t_r}\}_{r=1}^{s} = \{v_{(k+r)}\}_{r=1}^{v-1}$. Thus, the subsequence from the p^{th} to the t^{th} visit in F_C is a cyclic firing traversal of C with the form

$$v_k^{a_p} v_{(k+1)}^{a_{p+1}} v_{(k+2)}^{a_{p+2}} \ldots v_{(k+v-1)}^{a_{p+v-1}} v_k^{a_{p+v}}. \tag{14}$$

Now, we may assume that the firing of $v_{(k+i)}$ at the $(p+i)^{th}$ visit depends on the firing of $v_{(k+i-1)}$ at the $(p+i-1)^{th}$ visit and the subsequence from the p^{th} to the t^{th} visit in F_C cannot be reduced as in case 1. This follows from the hypothesis that F_C is not reducible. Furthermore, since v_k is assumed to be the first vertex of C which repeats in F_C, and all v vertices of C are present in the subsequence, it follows that $p = 1$. So, v_k must be the first vertex visited in the firing sequence F_C. Thus, $t = v + 1$ and F_C hs the form

$$v_k^{a_1} v_{(k+1)}^{a_2} v_{(k+2)}^{a_3} \ldots v_{(k+v-1)}^{a_v} v_k^{a_{v+1}} x_{v+2}^{a_{v+2}} \ldots x_m^{a_m}. \tag{15}$$

Specifically, case 2 must occur between nearest visits of any vertex which repeats in F_C; otherwise a reduction would be possible. Nearest visits of any vertex of C, in F_C, must necessarily be ν visits apart. From this observation and the fact that the first ν visits of F_C are cyclic, it follows that F_C must be cyclic.//

This result, along with the fact that we need only examine greedy firing sequences, allows us to immediately conclude that we may obtain a minimum scatter, legal firing sequence F_C for a disjoint directed circuit C in a marked directed graph G by examining all the greedy, cyclic, legal firing sequences of C which execute a given firing count vector Σ. There is exactly one such sequence starting at each legally firable vertex v_k of C at M_0. Since there can be at most ν vertices of C which are legally firable under a marking of C, this problem is at most ν times as complex as the problem of determining one cyclic firing sequence of C.

Specifically the greedy, cyclic firing sequence F_C^k, executing the firing count vector $\Sigma = [\sigma_0, \sigma_1, \ldots, \sigma_{\nu-1}]^T$, from an initial marking M_0 of C, is uniquely determined by its starting vertex v_k. Since we need only consider the case where σ_j is nonzero for each vertex v_j of C, it follows that F_C^k is at least ν visits long for each firable vertex v_k of C at M_0. The first ν visits of F_C^k are

$$v_k^{a_1}\, v_{(k+1)}^{a_2}\, v_{(k+2)}^{a_3} \ldots v_{(k+\nu-1)}^{a_\nu}, \tag{16}$$

where the a_j's are computed according to the recursive relation $a_j = \min\{\sigma_{(k+j-1)}, M_0(e_{(k+j-1)}) + a_{j-1}\}, j = 2, 3, \ldots \nu$, with $a_1 = M_0(e_k)$. Here, we have assumed the disabling number of vertex v_k is nonzero under the marking M_0. Otherwise, a zero scatter firing sequence exists starting at v_k, given by the above definition, where a_j is replaced with $\sigma_{(k+j-1)}$ for all $j \in \{1, 2, \ldots, \nu\}$. Let $\Sigma' = [\sigma_0', \sigma_1', \ldots, \sigma_{\nu-1}']^T$ be the firing count vector corresponding to the first ν visits of F_C^k and let $\Sigma'' = [\sigma_0'', \ldots, \sigma_{\nu-1}'']^T = \Sigma - \Sigma'$ be the residual firing count vector. Note that the firing count vector Σ' is some cyclic permutation of $[a_1, a_2, \ldots, a_\nu]^T$.

Theorem 6

If the firing numbers $\sigma_i, i = 0, 1, \ldots, \nu-1$ of a marked directed circuit C are expressed modulo the circuit token count Γ_C as $\sigma_i = s_i\Gamma_C + r_i$, where $s_i = \lfloor \sigma_i/\Gamma_C \rfloor$ and the remainder $r_i = \sigma_i \bmod (\Gamma_C)$, then

$$|s_i - s_j| \le 1, i, j \in \{0, 1, \ldots, \nu-1\}. \tag{17}$$

Proof

Consider any two firing numbers σ_i, σ_j of a marked directed circuit C with a circuit token count Γ_C. We make the following claim. Any two firing numbers σ_i, σ_j of a marked directed circuit C can differ by at most Γ_C. That is,

$$|\sigma_i - \sigma_j| \le \Gamma_C, i, j \in \{0, 1, \ldots, \nu-1\}. \tag{18}$$

If v_i is adjacent to v_j, then the difference $\sigma_i - \sigma_j$ is plus or minus the change in the marking on the edge connecting them. This change can be at most Γ_C in either direction. If v_i is not adjacent to v_j, then the difference $\sigma_i - \sigma_j$ can be expressed as the sum of the differential marking $M_* - M_0$ along the directed path in the circuit from v_i to v_j. Again, this difference can be at most, plus or minus Γ_C, thus establishing the claim.

Expressing the firing numbers σ_i, σ_j modulo Γ_C,

$$\sigma_i = s_i\Gamma_C + r_i, \quad \sigma_j = s_j\Gamma_C + r_j. \tag{19}$$

If $s_i \ne s_j$, then let $s_i > s_j$, without loss of generality. If $s_i \ne s_j + 1$, then

$$\sigma_i - \sigma_j = (s_i - s_j)\Gamma_C + r_i - r_j. \tag{20}$$

Since $s_i > s_j + 1$, $(s_i - s_j)\Gamma_C > 2\Gamma_C$. Clearly, we must have $|r_i - r_j| < \Gamma_C$ since $0 \le r_i, r_j \le \Gamma_C - 1$. It follows that

$$\sigma_i - \sigma_j = (s_i - s_j)\Gamma_C + r_i - r_j > \Gamma_C, \tag{21}$$

contradicting the claim established above.//

Theorem 7

The residual firing numbers have the property

$$\sigma_k'' \ge \sigma_{(k+1)}'' \ge \sigma_{(k+2)}'' \ge \ldots \ge \sigma_{(k+\nu-1)}'' \ge 0.$$

Proof

Assume, without loss of generality, that $k = 0$. Then, we must show that $\sigma_0'' \ge \sigma_1'' \ge \sigma_2'' \ldots \ge \sigma_{\nu-1}'' \ge 0$. The firing numbers σ_j satisfy the system

$$\sigma_{(j)} - \sigma_{(j+1)} = M_*(e_{(j+1)}) - M_0(e_{(j+1)}), j = 0, 1, \ldots, \nu-1. \tag{22}$$

The first $\nu-1$ of these equations form an independent set which can be written as

$$\sigma_{j+1} = \sigma_j + M_0(e_{j+1}) - M_*(e_{j+1}), j = 0, 1, \ldots, \nu-2. \tag{23}$$

If this system of equations is solved in terms of σ_0, the firing numbers are

$$\sigma_k = \sigma_0 + \sum_{j=1}^{k} M_0(e_j) - \sum_{j=1}^{k} M_*(e_j), k = 1, 2, \ldots, \nu-1. \tag{24}$$

Recursively applying the relation defining the a_j's,

$$a_k = \min\{\sigma_{k-1}, \sum_{j=0}^{k-1} M_0(e_j)\}, k = 1, 2, \ldots, \nu. \tag{25}$$

Since $\sigma_k' = a_{k+1}$, for $k = 0, 1, \ldots, \nu-1$,

$$\sigma_k'' = \sigma_k - \min\{\sigma_k, \sum_{j=0}^{k} M_0(e_j)\}, k = 0, 1, \ldots, \nu-1 \tag{26}$$

or,

$$\sigma_k'' = \max\{\sigma_k - \sum_{j=0}^{k} M_0(e_j), 0\}, k = 0, 1, \ldots, \nu-1. \tag{27}$$

Substituting the solution for σ_k in (24) into (27) gives

$$\sigma_k'' = \max\{\sigma_0 - M_0(e_0) - \sum_{j=1}^{k} M_*(e_j), 0\}, \quad k = 1, 2, \ldots, \nu-1. \tag{28}$$

Since the markings are nonnegative vectors, it follows that the numbers defined by

$$b_0 = \sigma_0 - M_0(e_0), \quad b_k = b_0 - \sum_{j=1}^{k} M_*(e_j), k = 1, 2, \ldots, \nu-1, \tag{29}$$

have the property $b_k \ge b_{k+1}$ for $k = 0, 1, \ldots, \nu-2$. Thus, the residual firing numbers $\sigma_k'' = \max\{b_k, 0\}$, for $k = 0, 1, \ldots, \nu-1$, must have the property

$$\sigma_0'' \ge \sigma_1'' \ge \sigma_2'' \ldots \ge \sigma_{\nu-1}'' \ge 0. \tag{30}$$

Since this result is independent of the vertex labeling, we must have

$$\sigma_k'' \ge \sigma_{(k+1)}'' \ge \sigma_{(k+2)}'' \ge \ldots \ge \sigma_{(k+\nu-1)}'' \ge 0, \tag{31}$$

for the residual firing numbers Σ'' corresponding to the firing sequence F_C^k.//

To completely characterize the greedy cyclic firing sequence F_C^k, two cases are to be considered.

Case 1: $\sigma''_{(k+v-1)} = 0$.
Here $v_{(k+v-1)}$ may be removed from C and Theorem 3 guarantees the existence of a zero scatter firing sequence executing Σ'' over the induced acyclic subgraph. The number of visits needed to execute Σ'' is just the number of nonzero entries in Σ''.

Case 2: $\sigma''_{(k+v-1)} \neq 0$.
In this case, all residual firing numbers are nonzero and the greedy cyclic firings of the vertices of C have placed all Γ_C tokens on edge e_k. Since any vertex of C can be fired at most Γ_C times per visit, each vertex v_i must be visited at least s_i times to execute σ_i'', where $s_i = \lfloor \sigma_i''/\Gamma_C \rfloor$. Since $s_{(k+v-1)}$ is a minimum s_i, then the entire circuit must be traversed at least $s_{(k+v-1)}$ times to execute Σ''. Each traversal will have the form

$$v_k^{\Gamma_C}\, v_{(k+1)}^{\Gamma_C}\, v_{(k+2)}^{\Gamma_C} \ldots v_{(k+v-1)}^{\Gamma_C}, \tag{32}$$

and the firing count Σ''' remaining after $s_{(k+v-1)}$ such traversals, is obtained by subtracting $s_{(k+v-1)}\Gamma_C$ from each entry in Σ''. If $s_{(k+v-1)}$ is zero, then $\Sigma''' = \Sigma''$ and such a traversal is not possible. Clearly, the residual firing numbers σ_i''' satisfy the property of Theorem 7. Therefore, $\sigma_k''' = (s_k - s_{(k+v-1)})\Gamma_C + r_k$ is a maximum σ_i'''. From Theorem 6 and Theorem 7, we have $s_k - s_{(k+v-1)} \leq 1$ and, therefore, vertex v_k need not be visited more than twice to execute σ_k'''. At this point, the number of visits needed to execute Σ''' may be counted. Specifically, two visits are required for each σ_i'' greater than Γ_C and one visit is required for each nonzero σ_i'' less than or equal to Γ_C. Thus, the number of visits needed to execute Σ''' is the number of nonzero entries in Σ''' plus the number of entries in Σ''' greater than Γ_C. Let w be the number of entries in Σ''' greater than Γ_C, where $0 \leq w \leq v-1$.

In either case, the firing sequence F_C^k may be formally written as

$$F_C^k = F_1^k F_2^k F_3^k \tag{33}$$

where

$$F_1^k = v_k^{\sigma'_k}\; v_{(k+1)}^{\sigma'_{(k+1)}}\, v_{(k+2)}^{\sigma'_{(k+2)}} \;\ldots\; v_{(k+v-1)}^{\sigma'_{(k+v-1)}},$$

$$F_2^k = [v_k^{\Gamma_C}\, v_{(k+1)}^{\Gamma_C}\, v_{(k+2)}^{\Gamma_C} \ldots v_{(k+v-1)}^{\Gamma_C}]^{s_{(k+v-1)}},$$

$$F_3^k = v_k^{\Gamma_C} \ldots v_{(k+w-1)}^{\Gamma_C}\, v_{(k+w)}^{r_{(k+w)}}\; v_{(k+w+1)}^{r_{(k+w+1)}} \;\ldots\; v_{(k+w+v)}^{r_{(k+w+v)}}, \tag{34}$$

and the visits $v_j^{\Gamma_C}$ are not present in F_3^k, if $w = 0$. The problem remains to determine k.

The above results are summarized in the following theorem, where

len (F_C^k) = length of F_C^k in visits,

$s_{min} = \min[\lfloor \sigma_i''/\Gamma_C \rfloor]$,

n_1 = number of nonzero entries in Σ'',

n_2 = number of nonzero entries in Σ''',

n_3 = number of nonzero entries greater than Γ_C in Σ'''.

Theorem 8
i) If $\sigma''_{(k+v-1)} = 0$, then $\mathrm{len}(F_C^k) = v + n_1$.
ii) If $\sigma''_{(k+v-1)} > 0$, then $\mathrm{len}(F_C^k) = (s_{min} + 1)v + n_2 + n_3$.//

The following algorithm determines a vertex v_k which should be fired first to arrive at a legal minimum scatter firing sequence F_C^k for a disjoint directed circuit C (which is fire-restricted only by the tokens on the circuit edges) in a marked directed graph G, given an

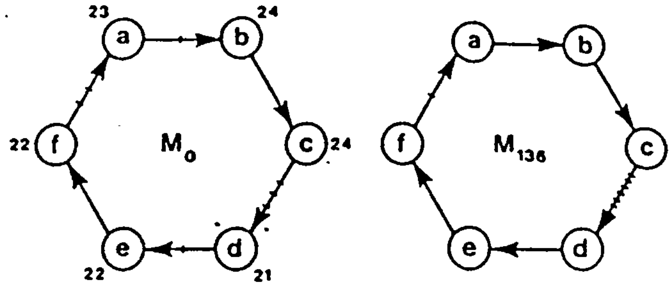

Figure 5: *Graph for illustrating Algorithm 1.*

executable firing count vector, from an initial marking M_0 of G. Correctness of the algorithm follows from Theorem 8. We assume the disabling numbers, and hence, the firing numbers are all nonzero. Let Σ be the firing count vector for the circuit vertices. Let C be labeled as in the previous subsection.

Algorithm 1
Step 1: For each legally firable vertex v_i of C at M_0, obtain a partition Σ'', Σ''' of the firing count vector Σ, defined by the system of equations:

$$\sigma'^{i}_{(i+j)\,mod(v)} = \min\left\{ \sum_{k=0}^{j} M_0\,(e_{(i+k)\,mod(v)}),\ \sigma_{(i+j)\,mod(v)} \right\},$$

$$\sigma''^{i}_{(i+j)\,mod(v)} = \sigma_{(i+j)\,mod(v)} - \sigma'^{i}_{(i+j)\,mod(v)}, \tag{35}$$

for $j = 0, 1, \ldots, v-1$. This step is of complexity $O(v^2)$.
Step 2: For each of the Σ'' vectors generated at step 1, the residual firing number $\sigma''^{i}_{(i-1)\,mod(v)}$ is a minimum one and σ_i'' is a maximum. Now, let $s_{min} = \lfloor \sigma''_{min}/\Gamma_C \rfloor$, where the residual count $\sigma''_{min} = \min[\sigma''^{i}_{(i-1)\,mod(v)}]$. Obtain all Σ''' vectors by reducing all firing numbers σ_j'' in each Σ'' by $s_{min}\Gamma_C$. If any of the Σ''' vectors contain zero entries, then proceed to step 3. Otherwise, proceed to step 4.
Step 3: From among all the Σ''' vectors, pick any one of them having a maximum number of zero entries. Let this vector be $\Sigma^{k'''}$. Stop.
Step 4: From among all the Σ''' vectors, pick any one of them having a maximum number of entries less than or equal to Γ_C. Let this vector be Σ_k'''. Stop.

At termination of this algorithm, the vertex v_k is the first vertex in the minimum scatter firing sequence of C as defined before. The minimum scatter firing sequence F_C^k can be determined from the vectors $\Sigma^{k''}$, $\Sigma^{k'''}$ and the number s_{min}, or by starting at vertex v_k and traversing the circuit in its direction, firing each vertex to its enabling number, and updating the enabling numbers, until all enabling numbers for the circuit are zero. It is easy to see that each step in the above algorithm is of complexity $O(v^2)$, and the overall complexity of this algorithm is $O(v^3)$.

Figure 5 is used to illustrate the steps involved in the above algorithm. It is assumed that this directed circuit C is a subgraph of some marked graph G, which is vertex-disjoint with all other directed circuits in G. The minimum firing count vector $\Sigma = [\sigma_a, \sigma_b, \sigma_c, \sigma_d, \sigma_e, \sigma_f]^T = [23, 24, 24, 21, 22, 22]^T$. There are four legally firable vertices of C at M_0. Step 1 produces the following partitions of Σ:

$$\Sigma^{a''} = [2, 3, 3, 6, 7, 7]^T,\ \Sigma^{a'''} = [21, 21, 21, 15, 15, 15]^T$$

$$\Sigma^{b''} = [7, 1, 1, 4, 5, 5]^T,\ \Sigma^{b'''} = [16, 23, 23, 17, 17, 17]^T$$

$$\Sigma^{d''} = [6, 7, 7, 3, 4, 4]^T,\ \Sigma^{d'''} = [17, 17, 17, 18, 18, 18]^T$$

$$\Sigma^{e''} = [3, 4, 4, 7, 1, 1]^T,\ \Sigma^{e'''} = [20, 20, 20, 14, 21, 21]^T$$

The minimum residual firing number is 14. Since $\Gamma_C = 7$, then $s_{min} = 2$. Thus, step 2 yields the vectors:

$$\Sigma^{a'''} = [7, 7, 7, 1, 1, 1]^T$$

$$\Sigma^{b'''} = [2, 9, 9, 3, 3, 3]^T$$

$$\Sigma^{d'''} = [3, 3, 3, 4, 4, 4]^T$$

$$\Sigma^{e'''} = [6, 6, 6, 0, 7, 7]^T$$

Since $\Sigma^{e'''}$ is the only residual vector with a zero entry, the greedy, cyclic firing sequence executing Σ, starting at vertex e, is a minimum scatter firing sequence leading from M_0 to M_{136}. From Theorem 8, the length of this sequence is 23 and the scatter is 17. The minimum firing sequence is

$$F_C = efa^3b^4c^4d^7(e^3f^7a^3b^7c^3d^7)^2e^3f^7a^4b^4c^4. \quad (36)$$

The above algorithm could be used to find a minimum scatter legal firing sequence for a marked directed graph G, in which all directed circuits are vertex-disjoint, as follows. First, identify all the directed circuits and condense them.[5] The resulting subgraph G' is acyclic. Let $v_1, v_2, \ldots, v_{v'}$ be the labels assigned to the vertices of G' according to a topological sort. Then, the sequence

$$F_G = F_1 F_2 F_3 \ldots F_{v'}, \quad (37)$$

is a minimum scatter firing sequence for G, where F_i is the minimum scatter subsequence, determined as before, if the vertex v_i corresponds to a condensed directed circuit. Otherwise, $F_i = v_i^{\sigma_i}$. Implementing the above procedure involves the following steps.

- Performing a depth-first search to determine all the disjoint-directed circuits. The complexity of this is $O(e + v)$,[5]
- Condensing the directed circuits and perform a topological sort on the resulting graph. This can also be achieved by means of a depth-first search.
- Applying Algorithm 1 to each directed circuit. This is of complexity $O(v^2)$.

The overall complexity of implementing the above procedure is $O(v^2)$.

On minimum scatter firing sequences for a general graph

We now outline an approach for determining minimum scatter firing sequences in the case of any general graph and then show that this problem is equivalent to the corresponding problem for the case of a strongly connected graph. Consider a marked directed graph $G = (V, E)$ where V is the vertex set and E is the edge set. Let $G_1 = (V_1, E_1), G_2 = (V_2, E_2) \ldots, G_k = (V_k, E_k)$ be the strongly connected components of G. Then, it is known that the sets $V_1\ V_2, \ldots, V_k$ form a partition of V.[5] Let G' be the graph obtained by contracting (short-circuiting) all the edges in each of the strongly connected components and removing the resulting self loops. Then, it is known that G' has no directed circuits. Let the vertex set of G' be denoted by $\{x_1, x_2, \ldots, x_k\}$. Assume that x_i corresponds to G_i and without loss of generality; let a topological sort of G' assign number i to vertex x_i. Then, it is easy to see that the sequence

$$F = F_1 F_2 \ldots F_k \quad (38)$$

is a minimum scatter firing sequence of G, where F_i is a minimum scatter firing sequence of G_i. This follows from the fact that after executing $F_1 F_2 \ldots F_{i-1}$, the vertices of G_i are fire-restricted only by the edges E_i of G_i. Note that the problem considered in the section discussing vertex-disjoint marked directed circuits is, in fact, the special case where each G_i is a directed circuit. The problem of determining a minimum scatter firing sequence for a general graph involves the following steps:

- determine all strongly connected components of G,[5]
- obtain the condensed graph G',
- perform a topological sort of the vertices of G'.
- find a minimum scatter firing sequence for each strongly connected component of G, and then
- a minimum scatter firing sequence F, of G, is given by

$$F = F_1 F_2 \ldots F_k, \quad (39)$$

where F_i is the minimum scatter subsequence of the strongly connected component corresponding to the i^{th} vertex in the topologically sorted vertex set of G'.

We now establish an upperbound on the enabling number μ_i of any vertex v_i under any marking reachable from a live initial marking M_0 on a strongly connected graph G. In a strongly connected directed graph G, each edge e_i belongs to a least one directed circuit. Since the circuit token count of all directed circuits of G must remain invariant under any legal sequence of vertex firings, the number of tokens on any edge e_i of a directed circuit can be no more than that circuit's token count. Extending this restriction to all directed circuits containing edge e_i implies that the token count on edge e_i cannot exceed the minimum circuit token count of all directed circuits containing edge e_i.

Let v_i be the terminal vertex of edge e_i. Now, every circuit containing edge e_i also contains vertex v_i. Thus, the enabling number μ_i of vertex v_i cannot exceed the minimum circuit token count of all directed circuits containing edge e_i. Applying this agrument to all edges encident into vertex v_i implies that the enabling number μ_i of any vertex v_i in a strongly connected graph G cannot exceed the minimum circuit token count of all directed circuits containing v_i. Let $\Gamma_{i,j}$ be the circuit token count of the j^{th} directed circuit containing vertex v_i, then

$$\mu_i \leq \min_j \{\Gamma_{i,j}\}, i \in \{1, 2, \ldots, v\}. \quad (40)$$

We may use this result to determine a lowerbound on the scatter of any firing sequence of a marked directed graph G. Since it is shown that only the strongly connected components of G need be considered, we present the result for a strongly connected graph. Let Γ_i denote the bound on the enabling number μ_i of vertex v_i, as defined above. Clearly, vertex v_i can be fired at most Γ_i times per visit. Thus, vertex v_i must be visited at least $\lceil \sigma_i / \Gamma_i \rceil$ times to execute σ_i. If v is the number of vertices with a nonzero firing count σ_i in a strongly connected graph G then for any legal firing sequence F executing Σ, we have

$$\text{scatter}(F) \geq \sum_{i=1}^{v} \lceil \sigma_i / \Gamma_i \rceil - v. \quad (41)$$

Summary

The following are the main contributions of this paper:

- An algorithmic proof of a theorem due to Murata is given.[2] The proof is based on an earlier work by Commoner, Holt, Even and Pnueli.[3]
- The concept of scatter in a firing sequence is introduced. Using the notion of a greedy firing policy, algorithms for generating minimum scatter firing sequences for different classes of graphs are presented.
- A general approach for determining minimum scatter firing sequences is discussed. It is shown that the problem in the general case is equivalent to the corresponding problem for the case of a strongly connected graph.

References

1. J.L. Peterson, "Petri Nets," *Computing Surveys*, Vol. 9, pp. 223–252, September 1977.
2. T. Murata, "Circuit Theoretic Analysis and Synthesis of Marked Graphs," *IEEE Trans. on Circuits and Systems*, Vol. CAS-24, No. 7, pp. 400–405, July 1977.
3. F.G. Commoner, A.N. Holt, S. Even, and A. Phueli, "Marked Directed Graphs," *J. of Computer and System Sciences*, Vol. 5, pp. 511–523, Oct. 1971.
4. M. Measures, P.A. Carr, B.D. Schriver, "A Distributed Operating System Kernel based on Data flow Principles," Proceedings of the 25th IEEE Computer Society International Conference, 1982.
5. M.N.S. Swamy and K. Thulasiraman, *Graphs, Networks and Algorithms*, Wiley-Interscience, New York, 1981.