Iterating the Branching Operation on a Directed Graph

Christos A. Athanasiadis*
DEPARTMENT OF MATHEMATICS
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
CAMBRIDGE, MA 02139, USA
E-mail: cathan@math.mit.edu

ABSTRACT

The branching operation D, defined by Propp, assigns to any directed graph G another directed graph $D(G)$ whose vertices are the oriented rooted spanning trees of the original graph G. We characterize the directed graphs G for which the sequence $\delta(G) = (G, D(G), D^2(G), \ldots)$ converges, meaning that it is eventually constant. As a corollary of the proof we get the following conjecture of Propp: for strongly connected directed graphs $G, \delta(G)$ converges if and only if $D^2(G) = D(G)$.

1. INTRODUCTION

Throughout this paper $G = (V, E)$ denotes a directed graph on a vertex set V, with multiple edges and loops allowed. All directed graphs we will consider are assumed to be finite. Let l be a nonnegative integer and recall that a path P in G is an alternating sequence $(u_0, e_1, u_1, \ldots, e_l, u_l)$ of vertices and edges of G such that each edge e_i has initial vertex u_{i-1} and terminal vertex u_i. We say that P is a path from u_0 to u_l and refer to e_1 as the *initial edge* of P if $l \geq 1$. An *oriented rooted spanning tree* on G, or simply a *rooted spanning tree* on G, is a spanning subgraph T of G together with a distinguished vertex r, called the root, such that for every $v \in V$ there is a unique path in T from v to r.

* Present address: Mathematical Sciences Research Institute, 1000 Centennial Drive, Berkeley, CA 94720. E-mail: athana@msri.org.

Journal of Graph Theory Vol. 24, No. 3, 257–265 (1997)

CCC 0364-9024/97/030257-09

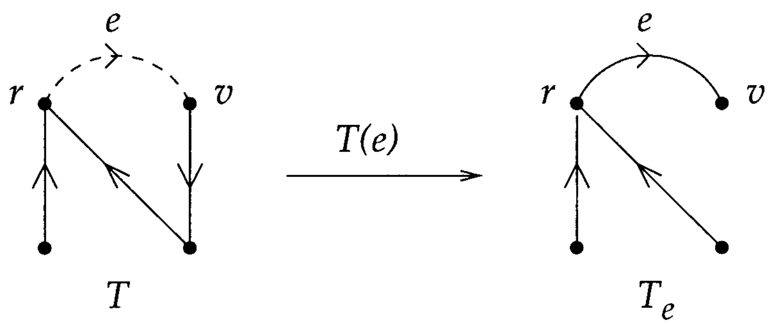

FIGURE 1. A re-rooting move.

We are concerned with a certain operation on directed graphs, denoted by D, which was defined by James Propp [5]. The object $D(G)$ is a directed graph whose vertex set is the set of rooted spanning trees of G. The edges in the new directed graph $D(G)$ are constructed as follows: Let T be a rooted spanning tree of G with root r and let $v \in V$. Given an edge $e \in E$ with initial vertex r and terminal vertex v, let T_e be the tree obtained from T by adding the edge e and deleting the edge with initial vertex v in $T \cup e$ (note that the root of T_e is v and that $T_e = T$ if e is a loop). Then add a directed edge in $D(G)$ from T to T_e. We denote this edge in $D(G)$ by $T(e)$. An example is illustrated in Figure 1.

We call D the *branching* operation. We use the letter D because it reminds us that $D(G)$ is some kind of a derived directed graph from G and since this was Propp's original notation. The idea of the construction of $D(G)$ appeared for the first time implicitly in the proof of the Markov chain-tree theorem by Anantharam and Tsoucas [1]. In this paper the authors needed to lift a random walk in G to a random walk in the set of arborescences of G, which coincides with the set of rooted spanning trees if G is strongly connected. On the other hand, Propp's motivation for defining $D(G)$ came from problems related to domino tilings of regions. Any domino tiling of a simply connected region can be obtained from any other tiling of the same region by a sequence of local changes, called ‘‘elementary moves’’ in [4]. Hence the set of domino tilings of such a region can be given the structure of a connected graph. An analogous result holds in the case of rooted spanning trees and the ‘‘re-rooting’’moves described in the definition of the branching operation. We state this as Lemma 2.2 in Section 2.

The directed graph $D(G)$ is a *covering space* of the directed graph G. This means that there exists a graph homomorphism p from $D(G)$ to G with the following property: if T and r are vertices of $D(G)$ and G respectively with $p(T) = r$ and if e is an edge in G with initial vertex r, then there exists a unique edge $\tilde{e}$ in $D(G)$ with initial vertex T such that $p(\tilde{e}) = e$. Indeed, to define such a map p we can simply map a rooted spanning tree T of G to its root r and the edge $T(e)$ of $D(G)$ with initial vertex T to the corresponding edge e in G with initial vertex r that gave rise to $T(e)$. The lift $\tilde{e}$ of e with initial vertex T is simply the edge $T(e)$.

This covering map seems to be what makes the branching operation interesting. It was crucial in showing that $D(G)$ has remarkable spectral properties [2] [3, Ch. 2]. More specifically, the eigenvalues of $D(G)$ can be computed directly from the eigenvalues of the induced subgraphs and the Laplacian matrix of G. Here we will be concerned with another aspect of the branching operation. For $i \geq 2$ let $D^i(G) = D(D^{i-1}(G))$ and consider the sequence

$$\delta(G) = (G, D(G), D^2(G), \ldots)$$

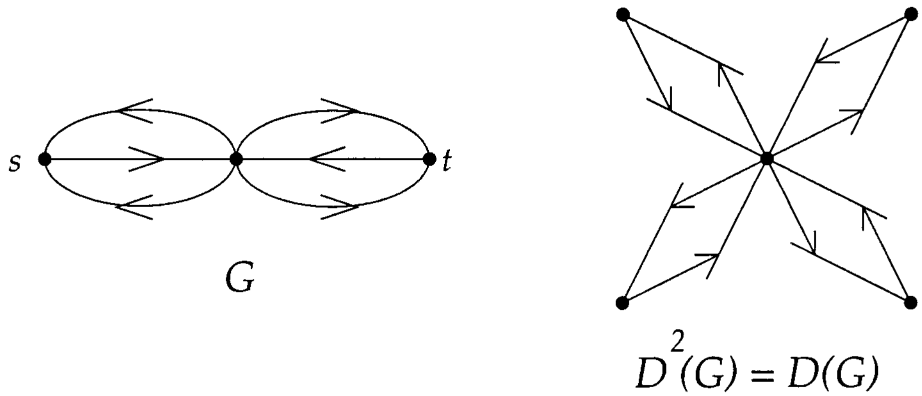

FIGURE 2. A counterexample to Propp's conjecture.

obtained by iterating D on G. Propp noted that $\delta(G)$ is very often a convergent (eventually constant) sequence, where equality of directed graphs is graph isomorphism. He asked for necessary and sufficient conditions on the graph G for $\delta(G)$ to be convergent and conjectured that, for strongly connected graphs $G, \delta(G)$ converges only when $D^2(G) = D(G)$ [5]. Moreover, he conjectured that this happens if and only if the simple paths in G, i.e., the paths in which each vertex appears at most once, satisfy a certain condition. Although his proposed condition turned out to be incorrect, we will show that a variant of this condition, given in the following definition, is true.

Definition 1.1. *We call G admissible if there are no two distinct vertices s, t in G having the following property: There exist two simple paths in G from s to t with different initial edges and similarly, two simple paths from t to s with different initial edges.*

In other words we require that for any two distinct vertices s, t of G, we have at most one choice for the initial edge of a simple path when moving from s to t or from t to s or in both directions. A different but equivalent formulation of the same condition is given as Lemma 2.1 and will be useful in Section 2.

Recall that G is said to be *strongly connected* if for any two distinct vertices s and t there exists a path in G from s to t. Our result, proved in the next two sections, can be stated as follows.

Theorem 1.2. *Let G be strongly connected. The sequence $\delta(G)$ converges if and only if G is admissible. Moreover, if $\delta(G)$ converges then $D^2(G) = D(G)$.*

Note that Definition 1.1 allows for the possibility that there are distinct simple paths in G from s to t and also from t to s as long as, in at least one direction, they all use the same initial edge. The example in Figure 2 shows that this extra freedom, which is what was missing from Propp's original formulation of the condition on s and t, is necessary for the theorem to be true.

The question of convergence of $\delta(G)$ for an arbitrary directed graph G is no more complicated. Either $D(G)$ is disconnected, and hence $\delta(G)$ stabilizes to the empty directed graph, or else $\delta(G)$ converges if and only if $\delta(H)$ does so, where $H = H(G)$ is a strongly connected induced subgraph of G, defined below. The details appear in Lemma 2.3.

More Notation and Definitions

We close this introduction with some basic constructions, notation and terminology which we will use in the proof of Theorem 1.2.

The *outdegree* of a vertex s in G, denoted $\mathrm{out}(s)$, is the number of nonloop edges of G emanating from s. Similarly we define the *indegree* of s and denote it by $\mathrm{in}(s)$. We denote by $r(G)$ the image of the vertex set of $D(G)$ under p, i.e., the set of roots of all rooted spanning trees of G. In other words, $r(G)$ is the set of vertices s of G such that for all $t \in V, t \neq s$, there exists a path in G from t to s. Note that $r(G) = V$ if and only if G is strongly connected. We denote by $H(G)$ the induced subgraph of G on $r(G)$, which is always strongly connected. We call G a *near monodromy* if for any two distinct vertices s and t there is at most one simple path from s to t.

We say that the path $P = (u_0, e_1, u_1, \ldots, e_l, u_l)$ in G visits a vertex u if $u = u_i$ for some $0 \leq i \leq l$. For any $0 \leq i \leq j \leq l$ we denote by $P(u_i, u_j)$ the part $(u_i, \ldots, e_j, u_j)$ of P which starts at u_i and ends at u_j. The *simplification* of P is a simple path from u_0 to u_l whose edges are some of the edges of P. To define it we start at u_0, choose the last edge e_k of P emanating from s and repeat the process with u_k until we reach u_l. Clearly the simplification of a simple path P is P itself and a path with one vertex and no edges if the initial and terminal vertex of P coincide.

If P is a path from s to t in G and Q is a path from t to u, then juxtaposition gives a well defined path PQ from s to u. The path PQ need not be simple even if P and Q are simple themselves. We denote the simplification of PQ by $P * Q$.

As a consequence of the covering property discussed above, we can lift paths in G uniquely up to $D(G)$, once the initial vertex is prescribed. In other words, if $p(T) = r$ then for any path Q in G with initial vertex r there is a unique path in $D(G)$, denoted by $T(Q)$, with initial vertex T such that $p(T(Q)) = Q$. Note that Q is not assumed to be simple. The notation $T(Q)$ emphasizes the dependence of the lifted path on both T and Q and is consistent with the notation $T(e)$, introduced earlier. In Section 3 we discuss how we can describe the endpoint of $T(Q)$ in terms of T and Q.

2. PROOF OF SUFFICIENCY

To prove the sufficiency part of Theorem 1.2 we need a few lemmas. The first lemma gives an alternative way to define admissibility.

Lemma 2.1. *G is admissible if and only if there is no ordered pair (s, t) of distinct vertices of G with the following property: There exist two simple paths in G from s to t with different initial edges and also two distinct simple paths from t to s.*

Proof. It suffices to show that the existence of an ordered pair (s, t) of distinct vertices with the property stated in the lemma implies the existence of two vertices as in Definition 1.1. Indeed, let $Q_1,\ Q_2$ be simple paths from s to t with different initial edges $e_1,\ e_2$ respectively and $P_1,\ P_2$ be distinct simple paths from t to s. Let u be the last vertex visited when walking P_1 or P_2 such that $P_1(t, u)$ and $P_2(t, u)$ coincide (so that $u = t$ if P_1 and P_2 have different initial edges). Since P_1 and P_2 are simple and distinct we have $u \neq s$. Then $P_1(u, s)$ and $P_2(u, s)$ are two simple paths in G from u to s with different initial edges. On the other hand the paths $Q_1 * P_1(t, u)$ and $Q_2 * P_1(t, u)$ give two simple paths in G from s to u with different initial edges, namely e_1 and e_2. Hence s and u have the property stated in Definition 1.1. ∎

The next lemma is from [2].

Lemma 2.2 ([2, Prop. 2.5] [3, Prop. 2.2.5]). *If G is strongly connected then so is $D(G)$.*

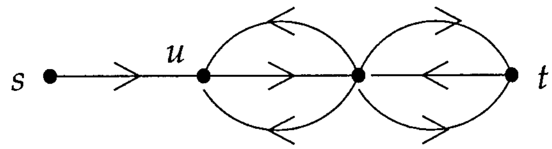

FIGURE 3. A counterexample to the stronger version of Lemma 2.4.

A directed graph G is strongly connected and a near monodromy if and only if for any two distinct vertices s and t there is a unique simple path in G from s to t. These directed graphs may be thought of as the directed analogues of undirected trees. A characterization is given in Section 4. The important property of strongly connected near monodromies for us is the fact that they satisfy the equation $D(G) = G$. Recall the definitions of $p, r(G)$ and $H(G)$ from the introduction.

Lemma 2.3. *If G is a near monodromy then $D(G)$ is isomorphic to the induced subgraph $H(G)$ of G. In particular, if G is a near monodromy and strongly connected then $D(G)$ is isomorphic to G. For any $G, D(G)$ is a disjoint union of copies of $D(H(G))$.*

Proof. Suppose G is a near monodromy. Then the map p gives the desired graph isomorphism. Indeed, since p is a covering map, it suffices to check that it induces a bijection between the vertex sets of $D(G)$ and $H(G)$. The vertex set $r(G)$ of $H(G)$ is by definition the image of the vertex set of $D(G)$ under p. Injectivity follows from the fact that G is a near monodromy, so that to any $r \in r(G)$ corresponds a unique rooted spanning tree of G with root r. Also $H(G) = G$ whenever G is strongly connected.

For the last statement note that if (u, v) is an edge of G and $u \in r(G)$ then $v \in r(G)$. The number of copies mentioned is the number of rooted spanning forests (disjoint unions of trees) of G with root set $r(G)$. ∎

The following two lemmas will also be essential in proving sufficiency.

Lemma 2.4. *Let G be admissible. Suppose s, t, u are distinct vertices of G and that there exist two simple paths from s to t with different initial edges and two distinct simple paths from t to u. Then there exist two distinct simple paths from s to u.*

Proof. Let P_1, P_2 be the two paths from s to t and Q_1, Q_2 be the two paths from t to u. If s is not visited by at least one of Q_1, Q_2, say Q_1, then $P_1 * Q_1$ and $P_2 * Q_1$ are simple paths from s to u. They have different initial edges, namely the initial edges of P_1 and P_2, respectively.

Suppose now that s is visited by both paths Q_1 and Q_2. Since G is admissible, by Lemma 2.1 the parts of Q_1 and Q_2 directed from t to s should coincide. But Q_1 and Q_2 are by assumption distinct, hence their parts $Q_1(s, u)$ and $Q_2(s, u)$ are distinct simple paths from s to u, as desired. ∎

The condition that the two paths from s to t have different initial edges cannot be dropped from the hypothesis. Figure 3 provides a counterexample to such a statement.

Lemma 2.5. *Let G be admissible. Then there exists a vertex s of G such that for any other vertex t of G there is at most one simple path in G from t to s. Moreover, if G is not a near monodromy, we can assume that the outdegree of s is at least 2.*

Proof. Suppose that such a vertex s (with outdegree ≥ 2) does not exist. Pick a vertex v_0 of G. If G is not a near monodromy we can choose v_0 so that $\mathrm{out}(v_0) \geq 2$. By assumption, there is a vertex $v_1 \neq v_0$ of G with the property that there are at least two distinct simple paths from v_1 to v_0, and we can certainly choose v_1 so that these two paths have different initial edges. Similarly, since clearly $\text{out}(v_1) \geq 2$, there is a vertex $v_2 \neq v_1$ and two simple paths from v_2 to v_1 with different initial edges. In this way we produce the sequence $v_0, v_1, v_2, \ldots$. Since our graph is finite, the sequence contains repeated vertices, say $v_i = v_j$ for some $i < j$. Clearly $i < j - 1$. By iterating Lemma 2.4 we see that there exist two distinct simple paths from v_{j-1} to v_i. By construction there exist two simple paths from $v_i = v_j$ to v_{j-1} with different initial edges. Lemma 2.1 implies that G is not admissible, a contradiction. ∎

We are now ready to state and prove the main result in this section.

Theorem 2.6. *If G is admissible then $D(G)$ is a near monodromy.*

Proof. We use induction on the number of nonloop edges $e(G)$ of G. The result is clear for $e(G) = 0$.

Suppose that $e(G) \geq 1$. First note that if G is a near monodromy then so is $D(G)$ since in this case, by Lemma 2.3, $D(G)$ is isomorphic to an induced subgraph of G.

Suppose now that G itself is not a near monodromy. Lemma 2.5 guarantees that there exists a vertex $s \in V$ with $\text{out}(s) \geq 2$, such that for every $t \in V, t \neq s$, there exists at most one simple path in G from t to s. If s is not in $r(G)$, i.e. there is no rooted spanning tree of G with root s, then we have $e(H(G)) < e(G)$. By Lemma 2.3, $D(G)$ is a disjoint union of copies of $D(H(G))$, hence the result follows by induction.

Lastly suppose that there exists a rooted spanning tree S of G with root s, in which case it is unique. Let $e_1, e_2, \ldots, e_k, k \geq 2$, be the nonloop edges emanating from s. Let $G_i, 1 \leq i \leq k$, be the directed graph obtained from G by deleting all edges e_j except e_i. Then each G_i is admissible with $e(G_i) < e(G)$. Hence, by induction, $D(G_i)$ is a near monodromy for each i. Let $\mathcal{T}_i$ be the set of vertices of $D(G_i)$ other than S. In other words, $\mathcal{T}_i$ is the set of rooted spanning trees of G having e_i as the edge emanating from s. Since any rooted spanning tree of G either equals S or has some e_i as the unique edge emanating from s, the sets $\mathcal{T}_i$ together with $\{S\}$ form a partition of the vertex set of $D(G)$. Note also that an edge in $D(G)$ with initial vertex in some $\mathcal{T}_i$ will have its terminal vertex in $\{S\} \cup \mathcal{T}_i$ and hence there is no edge in $D(G)$ directed from an element of $\mathcal{T}_i$ to an element of $\mathcal{T}_j$ for $i \neq j$. The fact that each $D(G_i)$ is a near monodromy now easily implies that $D(G)$ is a near monodromy as well. ∎

The next corollary follows from Theorem 2.6, Lemma 2.2 and Lemma 2.3 and proves one part of Theorem 1.2.

Corollary 2.7. *If G is strongly connected and admissible then $D^2(G) = D(G)$.*

3. PROOF OF NECESSITY

The other part of Theorem 1.2 states that if G is strongly connected but not admissible then $\delta(G)$ diverges. Note that for strongly connected directed graphs G, $D(G)$ has at least as many vertices as G, with equality if and only if G is a near monodromy. Thus Theorem 3.1, stated below, together with Lemma 2.2 implies that for G strongly connected but not admissible, the number of vertices of the graphs in $\delta(G)$ strictly increases. This completes the proof of Theorem 1.2.

One more piece of notation will be useful. Recall from the introduction that, given a path Q in G with initial vertex r and a rooted spanning tree T on G with root r, there is a unique path $T(Q)$ in $D(G)$ with initial vertex T which maps to Q under the homomorphism p, defined in Section 1. The path $T(Q)$ can be thought of as the way T gradually changes as we move along

Q. Walking an edge e of Q results in adding e to the tree currently visited by $T(Q)$, say S, and deleting the edge of $S \cup e$ with initial vertex the endpoint of e.

We denote by $T[Q]$ the terminal vertex of $T(Q)$. The tree $T[Q]$ can be described in terms of T and Q as follows: For any vertex u of G, the edge emanating from u in $T[Q]$ is the same as in T if Q never visits u and if it does, it is the edge of Q coming out of u the last time u is visited by Q. In particular, the endpoint of Q has no outcoming edge in $T[Q]$. We use this description to define $T[Q]$ for an arbitrary sequence of edges Q. The subgraph $T[Q]$ will not necessarily be a tree, but it will in all cases of interest. The crucial property of this construction for us is that $T[Q] = T[Q']$ if Q' is obtained from Q in one of the following ways:

(i) Deleting earlier occurrences of an edge.

(ii) Swapping two adjacent subsequences Q_1, Q_2 as long as any vertex v that is a vertex of some edge in both Q_1 and Q_2 has an emanating edge in Q that succeeds Q_1 and Q_2.

We are now ready to state and prove the theorem.

Theorem 3.1. *Suppose G is strongly connected. If G is not admissible then the same is true for $D(G)$.*

Proof. By assumption, there are two distinct vertices s and t in G with the following property: there exist two simple paths Q_1, Q_2 from s to t with different initial edges e_1, e_2 respectively and two simple paths P_1, P_2 from t to s with different initial edges. We want to show that the same is true for $D(G)$.

Let T be any rooted spanning tree on G with root s and let $T_1 = T[Q_2P_2Q_1P_1], T_2 = T[Q_1P_1Q_2P_2]$ be the final points of the paths $T(Q_2P_2Q_1P_1)$ and $T(Q_1P_1Q_2P_2)$ respectively. Note that T_1 and T_2 are distinct since the edges emanating from t in the two trees are different. We proceed to find pairs of paths in $D(G)$ from T_1 to T_2 and backwards, which will show that $D(G)$ is not admissible.

Let v be the first vertex other than t, visited by P_1, which is also visited by Q_1 (see Figure 4). Note that we could have $v = s$, but by assumption $v \neq t$. We construct two distinct paths R_1 and R_2 in $D(G)$ from T_1 to T_2 by setting

$$R_1 = T_1(Q_2P_2)$$

and

$$R_2 = T_1(Q_1P_2Q_1(s,v)P_1(v,s)Q_2P_2).$$

In the special case $v = s$ we have $R_2 = T_1(Q_1P_2Q_2P_2)$. Note that

$$T_1[Q_2P_2] = T[Q_2P_2Q_1P_1Q_2P_2] = T[Q_1P_1Q_2P_2] = T_2,$$

by cancelling the first occurrence of Q_2P_2 and that

$$\begin{aligned} T_1[Q_1P_2Q_1(s,v)P_1(v,s)Q_2P_2] &= T[Q_2P_2Q_1P_1Q_1P_2Q_1(s,v)P_1(v,s)Q_2P_2] \\ &= \text{(cancelling) } T[P_1(t,v)Q_1(v,t)Q_1(s,v)P_1(v,s)Q_2P_2] \\ &= \text{(permitted reordering) } T[P_1(t,v)Q_1P_1(v,s)Q_2P_2] \\ &= \text{(permitted reordering) } T[Q_1P_1Q_2P_2] = T_2. \end{aligned}$$

The way v was defined is essential for the next to last equality. It follows that R_1 and R_2 have indeed terminal vertex T_2.

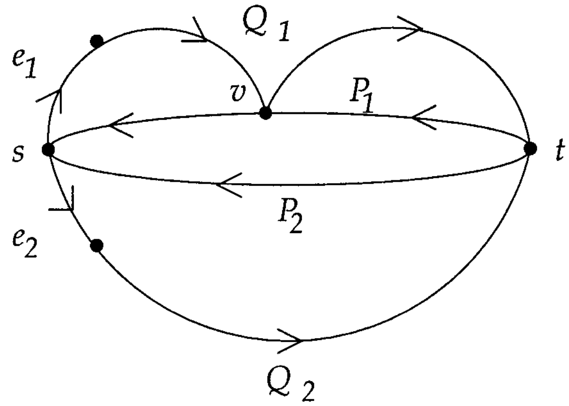

FIGURE 4. The four paths between s and t.

We now show that the simplifications of R_1 and R_2 (see Section 1) have different initial edges and therefore they provide us with the desired pair of paths from T_1 to T_2. Say that a tree not rooted at s is an e_1 *tree,* respectively e_2 *tree,* if the unique edge emanating from s in the tree is e_1, respectively e_2. We claim that the second vertex of the simplification of R_1 is an e_2 tree while the corresponding vertex for R_2 is e_1. For the first claim simply note that all trees visited by R_1, except for its endpoints T_1 and T_2, are e_2 trees. On the other hand, some of the trees not rooted at s and visited by R_2 are e_1 and some e_2. In fact, the trees visited by the first part $W = T_1(Q_1P_2Q_1(s,v)P_1(v,s))$ of R_2 not rooted at s are exactly the e_1 ones. Let S be the terminal vertex of W. It suffices to show that T_1 is not visited by $S(Q_2P_2)$, since then the second vertex of the simplification of R_2 will be a vertex of W and therefore e_1. The only possibility for T_1 to occur in $S(Q_2P_2)$ is that $T_1 = S$. But the edge emanating from t is the one travelled by P_1 in T_1 and the one travelled by P_2 in S since, by construction, t is not visited by $Q_1(s,v)P_1(v,s)$.

A completely symmetric argument gives the existence of analogous paths from T_2 to T_1 and shows that $D(G)$ is not admissible. ∎

4. ON DIRECTED ANALOGUES OF TREES

In this last section we give the characterization of strongly connected near monodromies promised in Section 2. We call such graphs simply *monodromies* and consider only loopless directed graphs in what follows. Thus a monodromy is a loopless directed graph $G = (V, E)$ having the following property: for any two distinct vertices $s, t \in V$ there exists a unique simple path in G from s to t. In particular, such a graph has no multiple edges. Equivalently, the monodromies are the strongly connected loopless directed graphs G satisfying $D(G) = G$.

Here we point out that we can construct all monodromies using a simple algorithmic procedure. To be more precise, suppose that G is a monodromy and that $q \in V$. Let's add to G a simple directed cycle with initial and terminal vertex q, i.e., a path

$$q = v_0 \to v_1 \to \cdots \to v_m = q,$$

where $m \geq 2$ is an integer and v_i, for $0 < i < m$, are new vertices added. It is not difficult to see that the new directed graph is also a monodromy. The next result says that any monodromy can be constructed from the loopless graph with one vertex, using a number of these operations.

Proposition 4.1. *Let $G = (V, E)$ be a monodromy. There exists a sequence of monodromies*

$$G_0 \subset G_1 \subset \cdots \subset G_n = G,$$

where G_0 is the loopless graph with one vertex, such that G_{i+1} can be obtained from G_i by adding a simple cycle, as described above.

Proof. The proof is by induction on the number of vertices of G, the result being clear if G has one or two vertices.

Let s be any vertex of G and let

$$s = u_0 \to u_1 \to \cdots \to u_r = t$$

be the longest simple path in G starting at s. Call this path P. We claim that out$(t) =$ in$(t) = 1$. To see this, note first that out(t), in $(t) \geq 1$ since by assumption, G is strongly connected with more than one vertex. Also if $(t, u) \in E$, then $u = u_i$ for some $0 \leq i \leq r - 1$ by maximality of P. So let $e_1 = (t, u_i) \in E$ for some $0 \leq i \leq r - 1$. If out$(t) \geq 2$ then $e_2 = (t, u_j)$ is another edge emanating from t for some $0 \leq j \leq r - 1$. Since G has no multiple edges, we may assume $i < j$. But then e_2 and e_1 followed by $P(u_i, u_j)$ give two distinct simple paths from t to u_j, contradicting the fact that G is a monodromy. Finally suppose that $(v, t) \in E$ for some $v \neq u_{r-1}$. By maximality of P, the unique simple path Q from s to v does not visit t, hence P and Q followed by (v, t) give two distinct simple paths from s to t and we again obtain a contradiction.

Let G' be the directed graph obtained from G by removing t and the two edges (u_{r-1}, t) and (t, u_i) coming into and out of t respectively, and adding an extra edge (u_{r-1}, u_i) if $i \neq r - 1$. G' is also a monodromy and has one vertex fewer than G. Hence the induction hypothesis applies and G' has the desired form. This easily implies the result for G as well. ∎

ACKNOWLEDGMENTS

I would like to thank James Propp for suggesting this problem and for his numerous helpful comments on this manuscript and an anonymous referee for suggesting several improvements on the presentation.

References

[1] V. Anantharam and P. Tsoucas, A proof of the Markov chain tree theorem, *Stat. Probab. Letters* **8** (1989), 189–192.

[2] C. A. Athanasiadis, Spectra of some interesting combinatorial matrices related to oriented spanning trees on a directed graph. *J. Alg. Combin.* **5** (1996), 5–11.

[3] C. A. Athanasiadis, Algebraic combinatorics of graph spectra, subspace arrangements and Tutte polynomials, Ph.D. thesis, MIT (1996).

[4] N. Elkies, G. Kuperberg, M. Larsen, and J. Propp, Alternating sign matrices and domino tilings, *J. Alg. Combin.* **1** (1992), 111–132.

[5] J. Propp, private communication.

Received September 20, 1994