Strongly Connected Orientations of Mixed Multigraphs

Fan R. K. Chung,* Michael R. Garey, and Robert E. Tarjan
AT&T Bell Laboratories, Murray Hill, New Jersey 07971

We study the problem of orienting all the undirected edges of a mixed multigraph so as to preserve reachability. Extending work by Robbins and by Boesch and Tindell, we develop a linear-time algorithm to test whether there is an orientation that preserves strong connectivity and to construct such an orientation whenever possible. This algorithm makes no attempt to minimize distances in the resulting directed graph, and indeed the maximum distance, for example, can blow up by a factor proportional to the number of vertices in the graph. Extending work by Chvátal and Thomassen, we then prove that, if a mixed multigraph of radius r has any strongly connected orientation, it must have an orientation of radius at most $4r^2 + 4r$. The proof gives a polynomial-time algorithm for constructing such an orientation.

1. INTRODUCTION

Let G be a mixed multigraph, i.e., a graph, possibly with multiple edges, each of which is either undirected or directed. We shall denote an edge with endpoints v and w by $\{v,w\}$ if it is undirected, $[v,w]$ if it is directed from v to w, and (v,w) if it is either undirected or directed from v to w. A *path* in G from vertex v_1 to vertex v_k is a sequence of edges $(v_1,v_2), (v_2,v_3), \ldots, (v_{k-1},v_k)$. The *length* of the path is $k - 1$. The path is a (simple) *cycle* if $v_1 = v_k$ and $v_1,v_2, \ldots, v_{k-1}$ are distinct. G is *strongly connected* if there is a path from v to w for every pair of vertices v,w. A *cut* $X,\overline{X}$ is a partition of the vertices of G into two nonempty parts, X and $\overline{X}$. An edge *crosses* the cut if it has one endpoint in each of X and $\overline{X}$. An edge is a *bridge* if it is the unique edge crossing some cut. An *orientation* of G is any multigraph formed by orienting all the undirected edges of G, i.e., converting each undirected edge $\{v,w\}$ into either $[v,w]$ or $[w,v]$. G is *orientable* if it has some orientation that is strongly connected. We are interested in the problem of testing orientability and finding strongly connected orientations.

Robbins [4] proved that an undirected graph is orientable if and only if it is connected

*Present address: Bell Communications Research, Morristown, New Jersey.

CCC 0028-3045/85/040477-08$4.00

and has no bridges. His proof yields a linear-time algorithm for constructing a strongly connected orientation if one exists: we carry out a depth-first search of the graph and orient every edge in the direction along which the search advances. (The same depth-first search can be used to test Robbins' condition; see [5].)

Boesch and Tindell [1] generalized Robbins' result to mixed multigraphs, showing that a mixed multigraph is orientable if and only if it is strongly connected and has no bridges. They claim that depth-first search can be used to find an appropriate orientation but provide no details. Our first result, presented in Section 2, is a linear-time algorithm for testing whether a mixed multigraph has a strongly connected orientation and finding one if it does. Our algorithm is a variant of Tarjan's algorithm [5] for finding strong components in a digraph using depth-first search.

The results of Robbins and of Boesch and Tindell say nothing about how much the distances between vertices can increase in the process of orienting edges while preserving strong connectivity. The paper of Chvátal and Thomassen [1] addresses this question in the case of undirected graphs. They prove that if an undirected graph of radius r is orientable, then it has an orientation of radius at most $r^2 + r$. Our second result, presented in Section 3, generalizes Chvátal and Thomassen's theorem to mixed multigraphs: We show that if a mixed multigraph of radius r is orientable, then it has an orientation of radius at most $4r^2 + 4r$. The proof is constructive and gives a polynomial-time algorithm for constructing such an orientation. Chvátal and Thomassen prove that their bound for radius blow-up of $r^2 + r$ is tight to within a constant factor in the worst case, which implies that our bound of $4r^2 + 4r$ is also tight to within a constant factor in the worst case. They also show that the problem of determining whether there is diameter-two orientation of an undirected graph is NP-complete, which implies that minimizing the radius or diameter of an orientation is NP-hard. Johnson and Pieroni [3] have proposed a linear programming approach to the related problem of orienting the edges in a graph with edge lengths so as to minimize the sum of the resulting distances between certain specified pairs of vertices.

2. FINDING A STRONGLY CONNECTED ORIENTATION IN LINEAR TIME

In order to devise an orientability-testing algorithm, we begin with two easy lemmas. Both follow from the orientability condition of Boesch and Tindell; we include proofs for completeness.

Lemma 1. If G is orientable, c is any cycle in G, and $\{v,w\}$ is any undirected edge on c, then the mixed multigraph formed by orienting $\{v,w\}$ in the direction consistent with the cycle is orientable.

Proof. Let $\vec{G}$ be a strongly connected orientation of G. Let $\vec{G}'$ be formed from $\vec{G}$ by reorienting every originally undirected edge on c in the direction consistent with c. If (x,y) is an edge of c, there is a path of directed edges in $\vec{G}'$ from x to y and another such path from y to x. (The two paths together form c.) It follows that $\vec{G}'$ is strongly connected, since for every path p in $\vec{G}$ there is a corresponding path p' in $\vec{G}'$, formed by substituting the appropriate path for each edge on p that is also on c. Since $\vec{G}'$ is an orientation of G, the lemma holds.

Lemma 2. Suppose G is orientable and $X,\bar{X}$ is any cut such that exactly one undirected edge, say, e, crosses $X,\bar{X}$, and all directed edges crossing the cut lead from X to $\bar{X}$. Then the graph G' formed by orienting e from $\bar{X}$ to X is orientable.

Proof. Immediate, since the only other choice, orienting e from X to $\bar{X}$, produces a graph that is not strongly connected. (No vertex in X is reachable from $\bar{X}$.)

Our orientability algorithm consists of two passes. The first pass orients all the undirected edges of G by applying Lemmas 1 and 2. The second pass checks that the resulting multigraph is strongly connected. To carry out the first pass, we perform a depth-first search of G, starting from an arbitrary start vertex. (In our discussion we shall assume some familiarity with the properties of depth-first search; see [5, 6].) We allow an advance to take place along an undirected edge in either direction, but along a directed edge only in the direction consistent with its orientation. We number the vertices from 1 to n as they are reached during the search; if not all vertices are reached, G is not orientable. To simplify the exposition, we shall assume that vertices are identified by number. The search partitions the edges of G into *tree edges,* those that when traversed during the search lead to previously unreached vertices, and *non-tree edges,* those that lead to previously reached vertices. The tree edges define a spanning tree rooted at the start vertex.

During the search, when we retreat along an undirected edge, we orient it. To guide the orientation process, we maintain, for each vertex v, a number $low(v)$ such that $1 \leq low(v) \leq v$. The *low* numbers function like the LOWPT numbers used to compute strong components by depth-first search [5]. It will always be the case that the current mixed multigraph contains a path directed from v to $low(v)$. We initialize $low(v) = v$ for each vertex v. We carry out the orientation process and the updating of *low* numbers by applying the appropriate case below whenever we retreat along an edge (v,w), backing up from w to v (see Fig. 1):

Case 1. (v,w) is a non-tree edge or $low(w) < w$. Replace $low(v)$ by $\min\{low(v), low(w)\}$. If (v,w) is undirected, orient it from v to w.

Case 2. (v,w) is a tree edge and $low(w) = w$. If (v,w) is directed, abort the search: G is not orientable. If (v,w) is undirected, orient it from w to v, and replace $low(w)$ by $low(v)$.

The entire first pass takes $O(n + m)$ time on a mixed multigraph of n vertices and m edges. The following lemmas summarize the important properties of this computation. Let $\vec{G}$ be the multidigraph formed from G by the first pass, if it does not abort.

Lemma 3. For any vertex v, $low(v) \leq v$, and there is a path from v to $low(v)$ in $\vec{G}$. If v is not the start vertex, then once the search retreats from v, $low(v) < v$. Thus if the search does not abort, there is path in $\vec{G}$ from any vertex v to 1, the start vertex.

Proof. All claims but the last in the lemma follow by induction on the number of retreats along edges performed during the search. the last claim follows immediately.

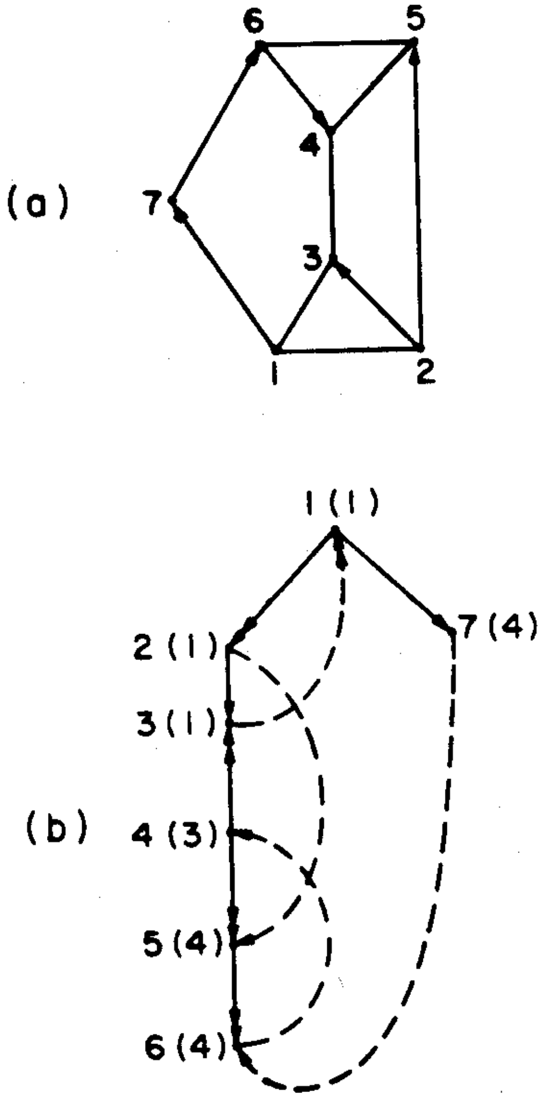

FIG. 1. Orienting a mixed multigraph. (a) Graph. Vertex numbers correspond to depth-first search order. (b) Spanning tree generated by depth-first search. Non-tree edges are dashed. Final *low* numbers are in parentheses. Undirected edges are directed as indicated by double arrows.

Lemma 4. If not all vertices are reached during the search, or if the search aborts, then G is not orientable. If G is orientable, then $\vec{G}$ is strongly connected.

Proof. Consider a retreat from a vertex w to a vertex v along a tree edge (v,w). We claim that $low(w) < w$ just before this retreat takes place if and only if there is an edge (x,y) in the original graph G such that x but not y is a descendant of w in the spanning tree generated by the search. (We regard w as a descendant of itself.) If there is such an edge, then the depth-first nature of the search implies $y < w \leq x$. After the search retreats along the edge (x,y), it will be the case that $low(x) \leq y$. This value will propagate back along the tree path from w to x, so that eventually $low(w) \leq y < w$. Conversely, if $low(w) < w$, $low(w)$ is not a descendant of w, and the path from w to $low(w)$ must contain an edge (x,y) such that x but not y is a descendant of w.

Suppose not all vertices are reached during the search. Then obviously G is not orientable, since not all vertices are reachable from the start vertex however the undirected edges are oriented. Suppose the search aborts in Case 2 while retreating along an edge (v,w), i.e., $low(w) = w$ just before the retreat. Let X contain the descendants of w and $\overline{X}$ the remaining vertices. Then, by the claim above every edge crossing the cut $X,\overline{X}$ is directed in G from $\overline{X}$ to X, which means that G is not orientable. This gives the first part of the lemma.

Suppose G is orientable. Consider a retreat along an undirected edge (v,w). If Case 2 occurs, then Lemma 2 and an argument like the one above imply that directing (v,w) from w to v preserves the orientability of G. Suppose Case 1 occurs. Lemma 3 implies that there is a path not containing (v,w) from w to an ancestor of v. By combining this with a path of tree edges we can form a cycle containing (v,w). By Lemma 1, directing (v,w) from v to w preserves the orientability of G. The second part of the lemma follows by induction on the number of edges oriented.

The second pass of the algorithm consists of a backward search from vertex 1 to test whether in $\vec{G}$ vertex 1 is reachable from every other vertex. If so, Lemma 3 implies that $\vec{G}$ is strongly connected and thus that G is orientable. If not, Lemma 4 implies that G is not orientable. The second pass takes $O(n + m)$ time.

Our algorithm gives an independent proof of Boesch and Tindell's necessary and sufficient condition for orientability, which we state here:

Theorem 1 [1]. A mixed multigraph is orientable if and only if it is strongly connected and has no bridges.

Proof. The necessity of the condition is obvious. To prove sufficiency, we consider the ways in which the algorithm can fail. If not all vertices are reached during the search of pass one, then G is not strongly connected. If the search aborts in Case 2, the proof of Lemma 4 implies that there is a cut $X,\overline{X}$ such that every edge in G that crosses the cut is directed from $\overline{X}$ to X, and hence G is not strongly connected.

Suppose G fails the test of the second pass. Let X be the set of vertices from which there is a path in $\vec{G}$ to 1, and let $\overline{X} \neq \emptyset$ contain the remaining vertices. All edges, if any, crossing the cut $X,\overline{X}$ are directed in $\vec{G}$ from X to $\overline{X}$. If there are no such edges, G is not strongly connected. If there is exactly one such edge, it is a bridge. Finally, suppose there are two or more such edges. Then all such edges must be directed in G, for neither Lemma 1 nor Lemma 2 permits orienting from X to $\overline{X}$ the last undirected edge crossing the cut. Thus in this case also G is not strongly connected. (There is no path from $\overline{X}$ to X.)

3. FINDING AN ORIENTATION OF SMALL RADIUS

The *radius* of a mixed multigraph G is the smallest integer r for which there is some vertex u (called a *center*) such that, for any vertex v, there are paths from u to v and from v to u of length at most r. (If G is not strongly connected, its radius is infinity.)

The orientation algorithm of Section 2 can cause a blow-up in radius from 1 to $n - 1$ if G has n vertices (consider a complete undirected graph). However, by choosing the orientation carefully, we can do much better. The following theorem generalizes a result of Chvátal and Thomassen [2], which applies only to undirected graphs but restricts the radius blow-up to $r^2 + r$:

Theorem 2. If G is a mixed multigraph of radius r, G has an orientation of radius at most $4r^2 + 4r$.

Our proof of Theorem 2 uses Chvátal and Thomassen's approach but is more complicated. We need a lemma:

Lemma 5. Let G be a mixed multigraph of radius r with center u. Then any edge incident to u is on a cycle of length at most $2r + 1$.

Proof. Consider a directed edge $[u,v]$. A shortest path from v to u has length at most r and does not contain $[u,v]$. Thus this path and $[u,v]$ form a cycle of length at most $r + 1$. The same argument applies to a directed edge $[v,u]$.

Consider an undirected edge $\{u,v\}$. Let X be the set of vertices w such that every shortest path from u to w contains $\{u,v\}$, and let Y be the set of vertices w such that every shortest path from w to u contains $\{u,v\}$. If $v \notin X$, there is another edge (u,v) in G and hence a cycle of length $2 \le r - 1$ containing $\{u,v\}$. If there is any vertex w in $X - Y$, there is a cycle of length at most $2r$ containing $\{u,v\}$, consisting of appropriate parts of a shortest path from u to w and a shortest path from w to u not containing $\{u,v\}$. The same argument applies if $Y - X \neq \emptyset$. Thus we can assume $X = Y \neq \emptyset$.

Let $\bar{X}$ be the set of vertices not in X. $\bar{X} \neq \emptyset$ since $u \in \bar{X}$. Since G contains no bridges, there must be an edge other than $\{u,v\}$, say, (x,y), crossing the cut $X,\bar{X}$. If $x \in X$ and $y \in \bar{X}$, there is a simple cycle of length at most $2r + 1$ containing $\{u,v\}$, formed from a shortest path from u to x, the edge (x,y), and a shortest path from y to u not containing $\{u,v\}$. The same argument applies if $x \in \bar{X}$ and $y \in X$, since $X = Y$. Thus the lemma is true in all cases.

We now proceed to prove Theorem 2 by induction on r. It holds trivially for $r = 0$. (In this case G consists of a single isolated vertex.) Suppose it holds for radius $r - 1$, and consider the case of radius r. Let u be a center of G. We shall show that there exists a vertex-induced subgraph H of G, containing u and all its incident edges, such that the undirected edges of H can be oriented to yield a multigraph $\vec{H}$ of radius at most $8r$. This will directly imply the theorem as follows: Form G' from G by condensing all vertices of H into a single vertex and deleting any resulting loops (edges with both endpoints the same). Since H contained u and all its incident edges, G' must have radius at most $r - 1$. Thus, by the induction hypothesis, G' has an orientation $\vec{G}'$ of radius at most $4(r - 1)^2 + 4(r - 1) = 4r^2 - 4r$. Expanding the condensed vertex back into the subgraph H and orienting its undirected edges according to their directions in $\vec{H}$, we obtain an orientation $\vec{G}$ of G with radius at most $4r^2 - 4r + 8r = 4r^2 + 4r$, as desired. The theorem will then follow the induction.

Therefore we need only show how to construct and orient an appropriate subgraph H. We do this in two stages. First we orient all undirected edges of G incident to u in such a way that each is contained in a cycle of length at most $4r$ in the resulting mixed multigraph. Then we construct and orient the remainder of H.

For the first stage, let $C_1, C_2, \ldots, C_k$ be a minimal collection of cycles, each of length at most $2r + 1$, containing every edge incident to u. (By "minimal" we mean that every cycle C_i contains some edge incident to u that is in no other cycle C_j.) We shall process these cycles in order by index, orienting their edges incident to u as we proceed and marking these edges as either *short* or *long*, so as to maintain the following invariants:

(i) Each marked edge is directed and in the current mixed multigraph is contained in a cycle of length at most $2r + 1$ if short, at most $4r$ if long;

(ii) Each unmarked edge is on some cycle C_i not yet processed (which need not be a cycle in the current mixed multigraph); and

(iii) no edge labeled long is on a cycle not yet processed.

We begin with C_1, simply orienting both of its edges incident to u in the direction consistent with C_1 (if they are not already so directed) and marking both of them short. Clearly (i)—(iii) hold, and the "short" cycle containing these edges will not be affected by orienting other edges incident to u.

In general, suppose C_i is the next cycle to be processed. By the minimality of C_1, $C_2, \ldots, C_k$, we know that C_i contains at least one edge incident to u that is unmarked. If C_i is still a cycle in the current mixed graph, we process C_i exactly as we did C_1; by (iii), the marked edge, if any, on C_i is already marked short, and remarking it short does not affect (i)–(iii).

If C_i is no longer a cycle, then exactly one of its two edges incident to u, say, edge e, is marked short and directed inconsistently with C_i, and the other edge incident to u, say, e', is unmarked and by minimality is on no cycle among $C_{i+1}, \ldots, C_k$. In this case we orient e' consistently with C_i and mark it long; it is on a cycle of length at most $4r$ consisting of part or all of C_i not including e and part or all of the short cycle containing e.

It is straightforward to verify that the general step preserves (i)–(iii). Hence, once all the cycles $C_1, C_2, \ldots, C_k$ have been processed, every edge incident to u will be directed and marked as either short or long, and thus in the resulting mixed multigraph each such edge is on a cycle of length at most $4r$.

For the second stage of the construction, grow a breadth-first spanning tree T_{out} out from u, and another breadth-first spanning tree T_{in} into u. For every vertex v such that $[u,v]$ is an edge, there is a path p_v in T_{in} from v to u of length at most $4r - 1$ not containing $[u,v]$. Similarly, for every vertex v such that $[v,u]$ is an edge, there is a path p'_v in T_{out} from u to v of length at most $4r - 1$ not containing $[v,u]$. Let X be the set of all vertices on some path p_v. For each vertex $v \notin X$ such that $[v,u]$ is an edge, let p''_v be the final part of p'_v from the last vertex in X to v. Let H be the subgraph of the current mixed multigraph induced by the set of vertices on some path p_v or p''_v.

Form $\vec{H}$ from H as follows. Orient every edge on a path p_v or p''_v in the direction consistent with the path; this creates no conflicts since all paths p_v are in T_{in} and all paths p''_v are in T_{out} and edge-disjoint with all paths p_v. Orient the remaining edges in H arbitrarily. Every edge $[u,v]$ is on a cycle in H of length at most $4r$ consisting of $[u,v]$ and p_v; every edge $[v,u]$ is on a cycle in H of length at most $8r$ consisting of (possibly) an edge $[u,w]$ and an initial part of p_w for some appropriate vertex w, followed by p''_v and $[v,u]$. (If p''_v consists of all of p'_v, the cycle consists only of p'_v and (v,u); if $v \in X$, the cycle consists only of $[u,w]$, all of p_w but the last edge, and $[v,u]$.) The subgraph $\vec{H}$ has radius at most $8r$ since u is a potential center satisfying that radius, and $\vec{H}$ includes all edges incident to u. Theorem 2 follows.

We close with two remarks. We can state Theorem 2 in terms of diameter instead of radius by noting that if r and d are the radius and diameter of a mixed multigraph, respectively, then $r \leq d \leq 2r$. Thus the construction used to prove Theorem 2 blows the diameter up from d to at most $8d^2 + 8d$. The proof of Theorem 2 leads to an algorithm for constructing an appropriate orientation that has a running time of $O(r^2 (n + m))$. This time bound is based on a naive implementation and may be improvable to $O(r(n + m))$.

References

[1] F. Boesch and R. Tindell, Robbins's theorem for mixed multigraphs. *Am. Math. Monthly* **87** (1980) 716–719.

[2] V. Chvátal and G. Thomassen, Distances in orientations of graphs. *J. Combinatorial Theory Ser. B* **24** (1978) 61–75.

[3] E. L. Johnson and P. Pieroni, A linear programming approach to the optimum network orientation problem. Presented at NETFLOW 83: International Workshop on Network Flow Optimization Theory and Practice, Pisa, Italy, 1983.

[4] H. E. Robbins, A theorem on graphs, with an application to a problem of traffic control. *Am. Math. Monthly* **46** (1939) 281–283.

[5] R. E. Tarjan, Depth-first search and linear graph algorithms. *SIAM J. Comput.* **1** (1972) 146–160.

[6] R. E. Tarjan, Finding dominators in directed graphs. *SIAM J. Comput.* **3** (1974) 62–89.

Received February, 1983
Accepted April, 1985