

Path-Based Depth-first Search for Strong and Biconnected Components

Harold N. Gabow *

February 28, 2000

Key words: Graph, depth-first search, strongly connected component, biconnected component, stack.

1 Introduction

Depth-first search, as developed by Tarjan and co-authors, is a fundamental technique of efficient algorithm design for graphs [23]. This note presents depth-first search algorithms for two basic problems, strong and biconnected components. Previous algorithms either compute auxiliary quantities based on the depth-first search tree (e.g., LOWPOINT values) or require two passes. We present one pass algorithms that only maintain a representation of the depth-first search path. This gives a simplified view of depth-first search without sacrificing efficiency.

In greater detail, most depth-first search algorithms (e.g., [23, 10, 11]) compute so-called LOWPOINT values that are defined in terms of the depth-first search tree. Because of the success of this method LOWPOINT values have become almost synonymous with depth-first search. LOWPOINT values are regarded as crucial in the strong and biconnected component algorithms, e.g. [14, pp. 94, 514]. Tarjan's LOWPOINT method for strong components is presented in texts [1, 7, 14, 16, 17, 21]. The strong component algorithm of Kosaraju and Sharir [22] is often viewed as conceptually simpler but it requires two passes over the graph. It is presented in texts [2, 4, 6, 25]. Tarjan's LOWPOINT biconnected component algorithm is presented in texts [1, 2, 4, 5, 7, 13, 14, 16, 17, 21, 25]. A two-pass biconnected component algorithm of Micali that avoids LOWPOINT values is sketched in [7, pp.67-68].

This paper presents strong and biconnected component algorithms that are based on the depth-first search path. This natural approach appears to have first been proposed by Purdom [19] and Munro [18] for strong components. It is regarded as requiring an extra data structure for set merging in order to be asymptotically efficient, and hence unlikely to be efficient in practice [23]. We present linear-time implementations of this approach for both strong and biconnected components. Our implementations use only stacks and arrays as data structures. A line-by-line pseudocode comparison of our algorithms with the tree-based algorithms of [23] shows the two approaches are similar in terms of lower level resource usage; performance differences are likely to be small or

*Department of Computer Science, University of Colorado at Boulder, Boulder, CO 80309. e-mail: hal@cs.colorado.edu

platform-dependent. Our algorithms show that the simpler path-based view of depth-first search suffices for these properties.

One can design other path-based depth-first search algorithms for properties such as ear decomposition [15], *st*-numbering [7], topological numbering, etc. The complete version of this paper [8] includes an algorithm to find the bridges of an undirected graph, leading to an immediate proof of Robbins' Theorem [20]. It also includes a simple articulation points algorithm, and a previously unpublished strong component algorithm of Tarjan that can be interpreted as path-based.

Section 2 presents our strong component algorithm and Section 3 presents the biconnected component algorithm. Appendix A proves a simple property of biconnected components. We conclude this section with some terminology.

Singleton sets are usually denoted by omitting set braces, e.g., for a set S and element x , $S - x$ denotes $S - \{x\}$. We assume all input graphs contain n vertices and m edges.

We use the following operations to manipulate a stack S : $\text{PUSH}(x, S)$ adds x to S at the (new) top of S . $\text{POP}(S)$ removes the value at the top of the stack and returns that value. $\text{TOP}(S)$ is the index of the value at the top of the stack. Hence $S[\text{TOP}(S)]$ is the value at the top of the stack.

2 Strong Components

Consider a digraph $G = (V, E)$. Two vertices are in the same *strong component* of G if and only if they are mutually reachable, i.e., there is a path from each vertex to the other. The *strong component graph* is formed by contracting the vertices of each strong component. Equivalently the strong component graph is the acyclic digraph, formed by contracting vertices of G , that has as many vertices as possible. In short we say the strong component graph is the finest acyclic contraction of G .

This characterization suggests the following high-level algorithm to find the strong component graph of $G = (V, E)$. See Fig. 1. The algorithm maintains a graph H that is a contraction of G with some vertices deleted. It also maintains a path P in H . Initially H is the given graph G .

If H has no vertices stop. Otherwise start a new path P by choosing a vertex v and setting $P = (v)$. Continue by growing P as follows.

To grow the path $P = (v_1, \dots, v_k)$ choose an edge (v_k, w) directed from the last vertex of P and do the following:

If $w \notin P$, add w to P , making it the new last vertex of P . Continue growing P .

If $w \in P$, say $w = v_i$, contract the cycle v_i, v_{i+1}, \dots, v_k , both in H and in P . P is now a path in the new graph H . Continue growing P .

If no edge leaves v_k , output v_k as a vertex of the strong component graph. Delete v_k from both H and P . If P is now nonempty continue growing P . Otherwise try to start a new path P .

It is easy to see that this algorithm forms the finest acyclic contraction of G . (For instance if no edge leaves v_k then v_k is a vertex of the finest acyclic contraction.) Thus the algorithm correctly computes the strong components.

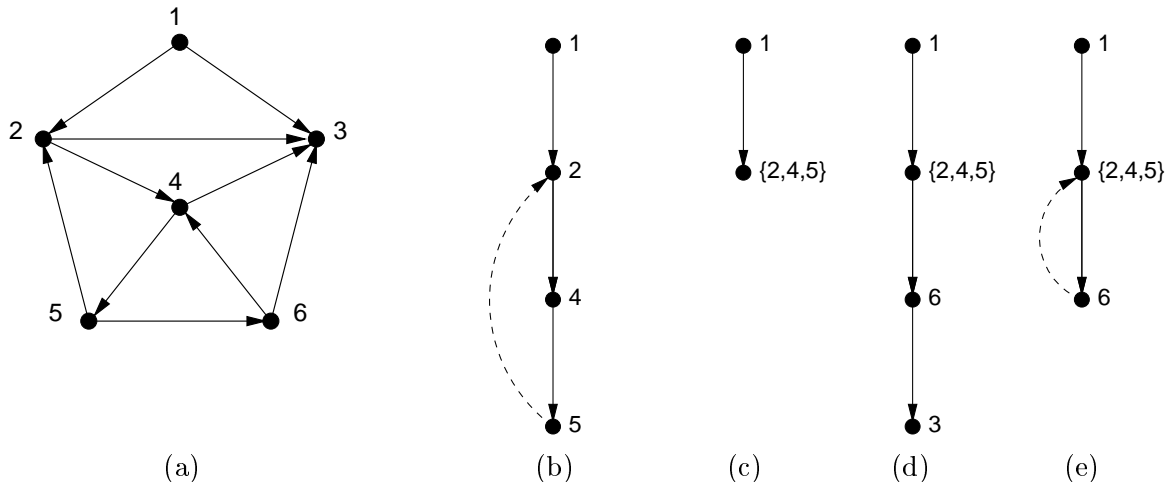


Figure 1: (a) Digraph G . (b)–(e) Path P (solid edges) in the first several steps of the algorithm. Strong component $\{3\}$ is output in (d).

This high-level algorithm was originally proposed by Purdom [19] and Munro [18]. The time for an efficient implementation is dominated by the time to keep track of the new vertices formed by contraction operations. Any data structure for disjoint set merging [6] can be used for this purpose. [19] and [18] use simple set-merging data structures, achieving total time $O(n^2)$ and $O(m + n \log n)$ respectively. Tarjan has shown set merging can be more efficiently, giving total time $O(m\alpha(m, n) + n)$ [24]. In fact the incremental tree set merging algorithm of [9] can be used. This reduces the time to $O(m + n)$, giving a linear time algorithm to find strong components. However the overhead of using incremental tree set merging may be significant in practice. Also the incremental tree algorithm requires a RAM machine and does not apply to a pointer machine. Now we give a simple list-based implementation that achieves linear time. The data structure is illustrated in Fig. 2.

Assume the vertices of the given graph G are numbered by consecutive integers from 1 to n . The algorithm numbers the strong components of G by consecutive integers starting at $n + 1$. It records the strong component number for each vertex (see (2) below).

Two stacks are used to represent the path P . Stack S contains the sequence of (original) vertices in P and stack B contains the boundaries between contracted vertices. More specifically S and B correspond to $P = (v_1, \dots, v_k)$ where $k = \text{TOP}(B)$ and for $i = 1, \dots, k$,

$$(1) \quad v_i = \{S[j] : B[i] \leq j < B[i + 1]\}.$$

When $k > 0$ we have $B[1] = 1$. Also when $i = k$ in (1) we interpret $B[k + 1]$ to be $\text{TOP}(S) + 1$.

An array $I[1..n]$ is used to store stack indices. It also stores the strong component number of a vertex when that number is known. More precisely for a given vertex v at any point in time,

$$(2) \quad I[v] = \begin{cases} 0 & \text{if } v \text{ has never been in } P; \\ j & \text{if } v \text{ is currently in } P \text{ and } S[j] = v; \\ c & \text{if the strong component containing } v \text{ has been deleted and numbered as } c. \end{cases}$$

Since there are only n vertices, there can be no confusion between an index j and a component number c in (2). A variable c is used to keep track of the component numbers.

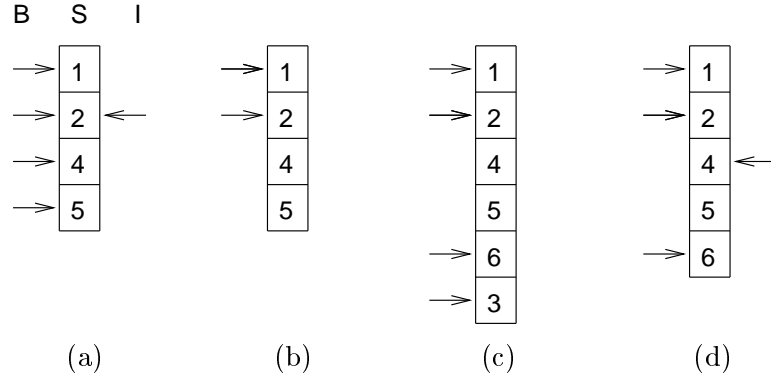


Figure 2: (a)–(d) show the data structure for Fig. 1(b)–(e) respectively. Stack S is shown. Arrows to the left of S represent stack B . Arrows to the right of S represent the entries of I that are used in contract steps. E.g. in (a) the algorithm reads $I[2] = 2$ and then contracts cycle 2, 4, 5 to get (b). In (c) $I[3]$ changes from 6 to 7, the latter being the strong component number of vertex 3.

The algorithm consists of a main routine **STRONG** and a recursive procedure **DFS**:

procedure **STRONG**(G)

1. empty stacks S and B ;
2. **for** $v \in V$ **do** $I[v] = 0$;
3. $c = n$;
4. **for** $v \in V$ **do** **if** $I[v] = 0$ **then** **DFS**(v);

procedure **DFS**(v)

1. **PUSH**(v, S); $I[v] = \text{TOP}(S)$; **PUSH**($I[v], B$); /* add v to the end of P */
2. **for** edges $(v, w) \in E$ **do**
3. **if** $I[w] = 0$ **then** **DFS**(w)
4. **else** /* contract if necessary */ **while** $I[w] < B[\text{TOP}(B)]$ **do** **POP**(B);
5. **if** $I[v] = B[\text{TOP}(B)]$ **then** { /* number vertices of the next strong component */
6. **POP**(B); increase c by 1;
7. **while** $I[v] \leq \text{TOP}(S)$ **do** $I[\text{POP}(S)] = c$ };

Theorem 2.1 *When **STRONG**(G) halts each vertex $v \in V$ belongs to the strong component numbered $I[v]$. The time and space are both $O(m + n)$.*

Proof: We will prove the first assertion of the theorem by showing that **STRONG** is a valid implementation of the high-level algorithm. We begin by specifying how the high-level algorithm will choose the edge (v_k, w) to grow P . Say that a vertex $w \in V$ becomes *active* (alternatively, *gets activated*) when it gets added to P as the new last vertex. The *most active* vertex is the currently active vertex that was activated most recently. To choose the next edge (v_k, w) let v be the most active vertex. Choose a previously unchosen edge directed from v , and use the corresponding edge of H as (v_k, w) . If all edges directed from v have been chosen then deactivate v . If this makes all vertices of v_k inactive then output v_k as the next strong component.

We must verify that this strategy correctly implements the high-level algorithm. This is easily done by verifying that v is a vertex of v_k , i.e., the most active vertex always belongs to the last vertex of P .

Now we prove that **STRONG** implements this version of the high-level algorithm. We assume that H , P and the deleted strong components are as specified by (1)–(2). The argument is by induction on the number of statements executed in **STRONG**. We will mention some points about the various statements and leave the remaining straightforward details of the induction to the reader. We refer to lines of pseudocode by the initial of the procedure name followed by the line number e.g., D7 is the last line of **DFS**.

When S4 is being executed, P is empty. (By convention the execution of a line or a statement excludes the execution of any recursive call.) During the execution of the loop of D2, v is the most active vertex.

In D3 if $0 < I[w] \leq n$ then D4 contracts a cycle or does nothing if (v, w) has already been contracted. If $I[w] > n$ then the component containing w has been deleted and D4 does nothing.

We turn to showing that the time and space are $O(m + n)$. We assume the given graph G is stored as a collection of adjacency lists. Observe that every vertex is pushed onto and popped from each stack S , B exactly once. Hence it is easy to see that the algorithm spends $O(1)$ time on each vertex or edge. \square

Comparing our code to the algorithm of [23], both methods use stack S ; our size n array I corresponds to a similar array that holds depth-first discovery numbers; our stack B corresponds to a size n array that holds **LOWLINK** values. Both S and B contain at most n entries at any time.

An algorithm almost identical to **STRONG** finds the bridges of an undirected graph. The high-level algorithm is based on the fact that contracting the vertices of a cycle does not change the bridges of a graph. The details are given in [8].

3 Biconnected Components

We present our algorithm for biconnected components in the language of hypergraphs. This is not logically necessary but it brings out the similarity to the strong components algorithm.

We start by reviewing basic definitions about hypergraphs [3, 15]. A *hypergraph* $H = (V, E)$ consists of a finite set V of *vertices* and a finite set E of *edges*, each edge a subset of V . A *path* is a sequence $(v_1, e_1, \dots, v_k, e_k)$ of distinct vertices v_i and distinct edges e_i , $1 \leq i \leq k$, with $v_1 \in e_1$ and $v_i \in e_{i-1} \cap e_i$ for every $1 < i \leq k$. The set of all vertices in edges of P is denoted $V(P) = \bigcup_{i=1}^k e_i$. A *cycle* is a path with the additional properties that $k > 1$ and $v_1 \in e_k$. A hypergraph is *acyclic* if it contains no cycle.

Notice that in a path P each vertex v_{i+1} , $1 \leq i < k$ belongs to $e_i - v_i$. For this reason the sets $e_i - v_i$ figure prominently in our algorithm (e.g. see (4) below). The algorithm also uses this operation on hypergraphs: To *merge* a collection of edges e_i , $i = 1, \dots, k$, add a new edge $\bigcup_{i=1}^k e_i$ and delete every edge of E contained in it (e.g., e_i). A *merging* of hypergraph H is a hypergraph formed by doing zero or more merges on H .

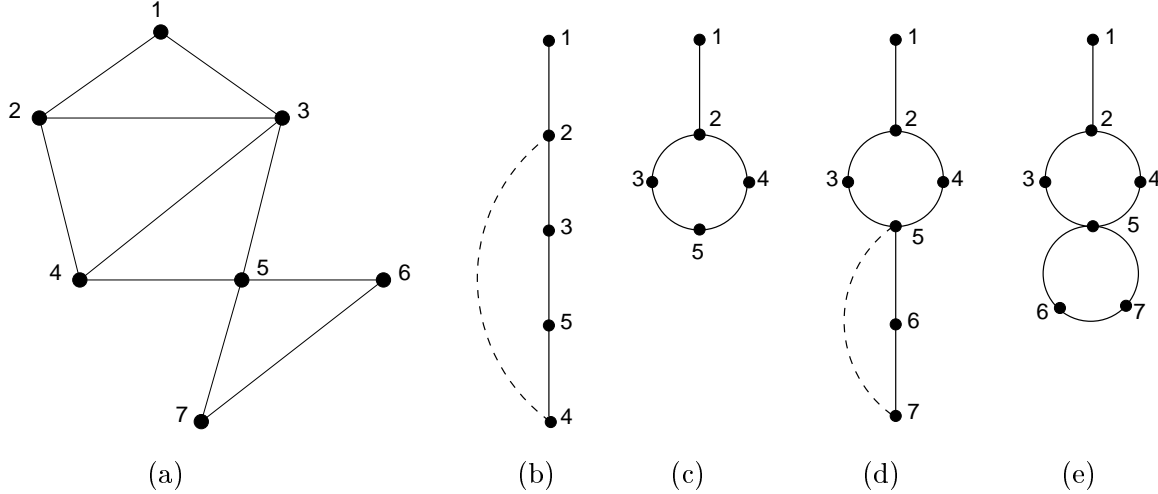


Figure 3: (a) Undirected graph G . (b)–(e) Path P (solid edges) in the first several steps of the algorithm. Biconnected component $\{5, 6, 7\}$ is output in (e).

Now consider an undirected graph $G = (V, E)$. Two distinct edges are in the same *biconnected component* of G if and only if some simple cycle contains both of them. This relation is easily seen to be an equivalence relation over the edges, so the biconnected components are well-defined. The “block-cutpoint tree” of a graph represents the biconnected components and cutpoints [12]. We will use a hypergraph variant of this notion: The *block hypergraph* H of G is the hypergraph formed by merging the edges of each biconnected component of G . H is an acyclic hypergraph. In fact H can be characterized as the finest acyclic merging of G , i.e., it is the acyclic hypergraph formed by merging edges of G that has as many hyperedges as possible. For completeness this characterization is proved in Appendix A.

The characterization suggests the following high-level algorithm to find the block hypergraph of $G = (V, E)$. See Fig. 3. The algorithm maintains a hypergraph H that is a merging of G with some edges deleted, and a path P in H . Initially H is the given graph G .

If H has no edges stop. Otherwise start a new path P by choosing an edge $\{v, w\}$ and setting $P = (v, \{v, w\})$ (choose the end v arbitrarily). Continue by growing P as follows.

To grow the path $P = (v_1, e_1, \dots, v_k, e_k)$ choose an edge $\{v, w\} \neq e_k$ with $v \in e_k - v_k$ and do the following:

If $w \notin V(P)$, add $v, \{v, w\}$ to the end of P . Continue growing P .

If $w \in V(P)$, say $w \in e_i - v_{i+1}$, merge the edges of the cycle $w, e_i, v_{i+1}, e_{i+1}, \dots, v_k, e_k, v, \{v, w\}$ to a new edge $e = \bigcup_{j=i}^k e_j$, both in H and in P . P is now a path ending with e (i.e., (v_i, e) has replaced $(v_i, e_i, \dots, v_k, e_k)$). Continue growing P .

If no edge leaves $e_k - v_k$, output e_k as an edge of the block hypergraph. Delete e_k from H and delete (v_k, e_k) from P . If P is now nonempty continue growing P . Otherwise try to start a new path P .

Correctness of this algorithm is based on two simple observations: When $v, \{v, w\}$ is added to P the result is a valid path, by the condition $v \in e_k - v_k$. When edges are merged they form a

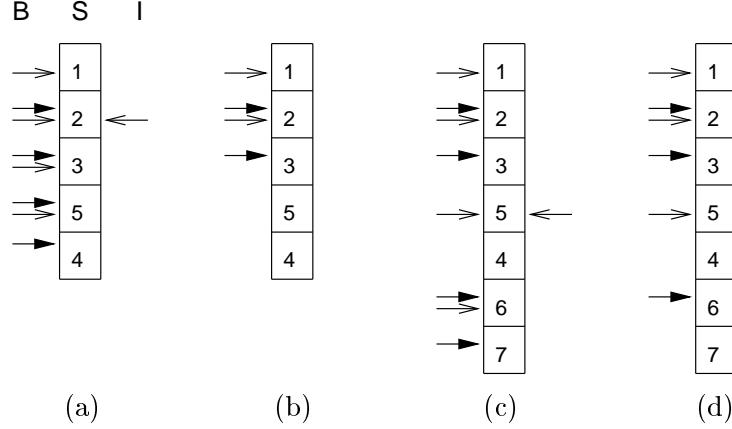


Figure 4: (a)–(d) illustrate the data structure for Fig. 3(b)–(e) respectively. S , B and I are represented as in Fig. 2. Every other arrowhead of B is drawn filled. E.g. in (a) the algorithm reads $I[2] = 2$ and then merges cycle 2, 3, 5, 4 to get (b). In (d) $I[6]$ and $I[7]$ change to 8, the number of the first biconnected component.

valid cycle, by the condition $\{v, w\} \neq e_k$. Now a straightforward inductive argument proves the algorithm correctly forms the finest acyclic merging of G , i.e., it finds the block hypergraph as desired.

As in Section 2 we give a list-based implementation that achieves linear time. The data structure is illustrated in Fig. 4. As before assume the vertices of G are numbered by consecutive integers from 1 to n . The algorithm numbers the biconnected components of G by consecutive integers starting at $n + 1$. The biconnected components are represented by assigning a number $I[v]$ to each vertex v in such a way that each edge $\{v, w\}$ belongs to the biconnected component with number $\min\{I[v], I[w]\}$ (see (5) below).

Two stacks are used to represent the path P . Stack S contains the vertices $V(P)$ and stack B represents the boundaries between edges of P , two vertices per boundary. More specifically S and B correspond to $P = (v_1, e_1, \dots, v_k, e_k)$, where $\text{TOP}(B) = 2k$ and for $i = 1, \dots, k$,

$$(3) \quad v_i = S[B[2i - 1]]; \quad (3)$$

$$(4) \quad e_i - v_i = \{S[j] : B[2i] \leq j < B[2i + 2]\}. \quad (4)$$

Thus in Fig. 4 the open arrows of B point to the vertices v_i of P . The filled arrows demarcate the sets $e_i - v_i$; these sets are the “nonfirst” vertices of edges e_i of P . When P is nonempty we have $B[i] = i$ for $i = 1, 2$. Also when $i = k$ in (4) we interpret $B[2k + 2]$ to be $\text{TOP}(S) + 1$.

As in the strong components algorithm an array I stores stack indices as well as biconnected component numbers. More precisely for a given vertex v at any point in time,

$$(5) \quad I[v] = \begin{cases} 0 & \text{if } v \text{ has never been in } P; \\ j & \text{if } v \text{ is currently in } P \text{ and } S[j] = v; \\ c & \text{if the last biconnected component containing } v \text{ has been output} \\ & \text{and numbered as } c. \end{cases} \quad (5)$$

As before there can be no confusion between an index j and a component number c in (5). A variable c is used to keep track of the component numbers.

The algorithm consists of a main routine BICONN and a recursive procedure DFS:

procedure BICONN(G)

1. empty stacks S and B ;
2. **for** $v \in V$ **do** $I[v] = 0$;
3. $c = n$;
4. **for** $v \in V$ **do** **if** $I[v] = 0$ and v is not isolated **then** DFS(v);

procedure DFS(v)

1. PUSH(v, S); $I[v] = \text{TOP}(S)$; **if** $I[v] > 1$ **then** PUSH($I[v], B$); /* create a filled arrow on B */
2. **for** edges $\{v, w\} \in E$ **do**
3. **if** $I[w] = 0$ **then** { PUSH($I[v], B$); DFS(w) /* create an open arrow on B */ }
4. **else** /* possible merge */ **while** $I[v] > 1$ and $I[w] < B[\text{TOP}(B) - 1]$ **do** {POP(B); POP(B)};
5. **if** $I[v] = 1$ **then** $I[\text{POP}(S)] = c$
6. **else if** $I[v] = B[\text{TOP}(B)]$ **then** {
7. POP(B); POP(B); increase c by 1;
8. **while** $I[v] \leq \text{TOP}(S)$ **do** $I[\text{POP}(S)] = c$ };

In many situations line B4 can be simplified: If G is known to be a connected graph B4 can be replaced by a single call DFS(v) (for any vertex v). If G has no isolated vertices, i.e., every vertex is on at least one edge, the second part of the **if** test of B4 can be dropped. Also moving the code for DFS for the case $I[v] = 1$ into B4 allows DFS itself to be simplified.

Theorem 3.1 *When BICONN(G) halts any edge $\{v, w\} \in E$ belongs to the biconnected component numbered $\min\{I[v], I[w]\}$. The time and space are both $O(m + n)$.*

Proof: The argument is similar to Theorem 2.1 and uses the conventions introduced in that proof. We prove the first assertion of the theorem by showing that BICONN is a valid implementation of the high-level algorithm.

We first specify how the high-level algorithm chooses the pair $v, \{v, w\}$ to grow P . Say a vertex $w \in V$ becomes *active* the first time it gets added to P . As before the *most active* vertex is the currently active vertex that was activated most recently. To choose $v, \{v, w\}$, let v be the most active vertex. Choose a previously unchosen edge $\{v, w\}$. If all edges incident to v have been chosen then deactivate v . If P is nonempty and this makes all vertices of $e_k - v_k$ inactive then output e_k as the next edge of the block hypergraph.

Note that it is possible to have P empty and a vertex v active. This can occur if v was the previous first vertex of P . In this case the above strategy starts a new path $P = (v, \{v, w\})$ by adding an edge incident to v . On the other hand it is possible to have P empty and no vertex v active. In this case when a new path $P = (v, \{v, w\})$ is started, by convention v becomes active before w . This convention ensures that the most active vertex is always unique.

This strategy correctly implements the high-level algorithm. To prove this we need only check that when P is nonempty $v \in e_k - v_k$, for v the most active vertex.

We will use another property of the implementation: When it chooses the pair $v, \{v, w\}$ to grow P , if $w \in V(P) - e_k$ then w is currently active. To show this note that $w \in V(P)$ implies w has been activated. Also $w \notin e_k$ implies $w, \{w, v\}$ has not been chosen to grow P (since after an edge is chosen, its ends belong to a common edge of P). This implies w is still active.

Now we prove that **BICONN** implements the above version of the high-level algorithm. We assume that H, P and the deleted biconnected components are as specified by (3)–(5). The argument is by induction on the number of statements executed in **BICONN**. We only mention the most important points about the various statements, leaving the remaining details of the induction to the reader.

When **B4** is being executed, P is empty. During the execution of the loop of **D2**, v is the most active vertex. (Note that if $I[v] = 1$ then P is empty during the execution of **D2**. In this case we also have $\text{TOP}(S) = 1$ and $\text{TOP}(B) = 0$.)

In **D3** suppose $I[w] > n$. Then the last biconnected component containing w has been deleted and **D4** does nothing. Suppose $0 < I[w] \leq n$. If $w \in e_k$ then $I[w] \geq I[v_k] = B[2k - 1]$ (by (3)) so **D4** does nothing. In the remaining case choose index i so $w \in e_i - v_{i+1}$, $1 \leq i < k$. As noted for the high-level algorithm, w is an active vertex in $V(P)$. Our choice rule implies v_{i+1} is the most recently activated vertex of e_i that is still active. Thus $I[v_{i+1}] > I[w] \geq I[v_i]$. Equivalently by (3), $B[2i + 1] > I[w] \geq B[2i - 1]$. This shows **D4** merges the same cycle as the high-level algorithm.

The test of **D6** checks whether or not the last edge e_k of P consists of v and its successors on S , plus vertex v_k . Hence **D8** labels vertices according to (5).

For the time and space bounds observe that every nonisolated vertex is pushed onto S exactly once. It is also pushed onto an even entry of B at most once. Hence it is easy to see that the algorithm spends $O(1)$ time on each vertex or edge. \square

Comparing our code to the algorithm of [23], our stack S (which has at most n entries) corresponds to a stack of edges (which has at most m entries). Our size n array I corresponds to a similar array that holds depth-first discovery numbers. Our stack B (which has at most $2n$ entries) corresponds to a size n array that holds **LOWPT** values.

Acknowledgments

We thank San Skulrattanakulchai for helpful suggestions.

Appendix A Characterization of the Block Hypergraph

Lemma A.1 *The block hypergraph of a graph G is the finest acyclic merging of G .*

Proof: We first show the block hypergraph H is acyclic. A biconnected component of G is a connected subgraph of G . Hence a cycle in H gives a cycle in G that contains edges from at least two distinct biconnected components. This is impossible.

To show H is the finest acyclic merging let K be an acyclic merging of G . Any cycle of G is contained entirely in one edge of K . Thus any biconnected component is contained in one edge of K . \square

References

- [1] A.V. Aho, J.E. Hopcroft and J.D. Ullman, *The Design and Analysis of Computer Algorithms*, Addison-Wesley, Reading MA, 1974.
- [2] A.V. Aho, J.E. Hopcroft and J.D. Ullman, *Data Structures and Algorithms*, Addison-Wesley, Reading MA, 1983.
- [3] C. Berge, *Hypergraphs: Combinatorics of Finite Sets*, North-Holland, NY, 1989.
- [4] G. Brassard and P. Bratley, *Algorithmics: Theory & Practice*, Prentice-Hall, Englewood Cliffs NJ, 1988.
- [5] G. Brassard and P. Bratley, *Fundamentals of Algorithmics*, Prentice-Hall, Englewood Cliffs NJ, 1996.
- [6] T.H. Cormen, C.E. Leiserson and R.L. Rivest, *Introduction to Algorithms*, McGraw-Hill, NY, 1990.
- [7] S. Even, *Graph Algorithms*, Computer Science Press, Potomac, MD, 1979.
- [8] H.N. Gabow, Path-based depth-first search for strong and biconnected components, Tech. Rept. CU-CS-890-99, revised version, Dept. of Computer Science, University of Colorado at Boulder, 2000.
- [9] H.N. Gabow and R.E. Tarjan, “A linear-time algorithm for a special case of disjoint set union,” *J. Comp. and System Sci.*, 30, 2, 1985, pp. 209–221.
- [10] J.E. Hopcroft and R.E. Tarjan, “Dividing a graph into triconnected components,” *SIAM J. Comput.*, 2, 1973, pp. 135–158.
- [11] J.E. Hopcroft and R.E. Tarjan, “Efficient planarity testing,” *J. ACM*, 21, 4, 1974, pp. 549–568.
- [12] F. Harary, *Graph Theory*, Addison-Wesley, Reading MA, 1969.
- [13] E. Horowitz, S. Sahni, S. Rajasekaran, *Computer Algorithms*, Computer-Science Press, NY, 1998.
- [14] D.E. Knuth, *The Stanford Graphbase: A Platform for Combinatorial Computing*, Addison-Wesley, Reading MA, 1993.
- [15] L. Lovász, *Combinatorial Problems and Exercises*, 2nd Edition, North-Holland, NY, 1993.
- [16] U. Manber, *Introduction to Algorithms: A Creative Approach*, Addison-Wesley, Reading MA, 1989.
- [17] K. Melhorn, *Data Structures and Algorithms 2: Graph Algorithms and NP-Completeness*, Springer-Verlag, NY, 1984.

- [18] I. Munro, Efficient determination of the strongly connected components and transitive closure of a directed graph, Department of Computer Science, University of Toronto, 1971.
- [19] P.W. Purdom, A transitive closure algorithm, Tech. Rept. 33, Computer Sciences Department, University of Wisconsin, Madison, 1968.
- [20] H.E. Robbins, “A theorem on graphs with an application to a problem of traffic control,” *American Math. Monthly*, 46, 1939, pp. 281–283.
- [21] R. Sedgwick, *Algorithms in C*, Addison-Wesley, Reading MA, 1990.
- [22] M. Sharir, “A strong-connectivity algorithm and its application in data flow analysis,” *Computers and Mathematics with Applications* 7, 1, 1981, pp. 67–72.
- [23] R.E. Tarjan, “Depth-first search and linear graph algorithms,” *SIAM J. Comput.*, 1, 2, 1972, pp. 146–160.
- [24] R.E. Tarjan, “Efficiency of a good but not linear set union algorithm,” *J. ACM*, 22, 2, 1975, pp. 215–225.
- [25] M.A. Weiss, *Data Structures and Algorithm Analysis in C++*, Addison-Wesley, Reading MA, 1999.