

# On the logspace shortest path problem

Boris Brimkov and Illya V. Hicks

Computational & Applied Mathematics, Rice University, Houston, TX 77005 USA  
boris.brimkov@rice.edu, ivhicks@rice.edu

**Abstract.** In this paper, we reduce the logspace shortest path problem to biconnected graphs; in particular, we present a logspace shortest path algorithm for general graphs which uses a logspace shortest path oracle for biconnected graphs. We also present a linear time logspace shortest path algorithm for graphs with bounded vertex degree and biconnected component size, which does not rely on an oracle. The asymptotic time-space product of this algorithm is the best possible among all shortest path algorithms.

**Keywords:** logspace algorithm, shortest path, biconnected graph, bounded degree

## 1 Introduction

The logspace computational model entails algorithms which use a read-only input array and  $O(\log n)$  working memory. For general graphs, there is no known deterministic logspace algorithm for the shortest path problem. In fact, the shortest path problem is NL-complete, so the existence of a logspace algorithm would imply that  $L=NL$  [9]. In this paper, we reduce the logspace shortest path problem to biconnected graphs, and present a linear time logspace shortest path algorithm for parameter-constrained graphs.

An important result under the logspace computational model which is used in the sequel is Reingold's deterministic polynomial time algorithm [13] for the undirected  $st$ -connectivity problem (USTCON) of determining whether two vertices in an undirected graph belong to the same connected component. There are a number of randomized logspace algorithms for USTCON (see, for example, [2, 6, 11]) which perform faster than Reingold's algorithm but whose output may be incorrect with a certain probability.

There are also a number of logspace algorithms for the shortest path problem and other graph problems on special types of graphs (see [1, 4, 5, 9, 12]). As a rule, due to time-space trade-off, improved space-efficiency is achieved on the account of higher time-complexity. Often the trade-off is rather large, yielding time complexities of  $O(n^c)$  "for some constant  $c$  significantly larger than 1" [10]. In particular, the time complexity of Reingold's USTCON algorithm remains largely uncharted but is possibly of very high order. The linear time logspace shortest path algorithm presented in this paper avoids this shortcoming, at the expense of some loss of generality. In fact, its time (and space) complexity is the

best possible, since a hypothetical sublinear-time algorithm would fail to print a shortest path of length  $\Omega(n)$ .

This paper is organized as follows. In the next section, we recall some basic definitions and introduce a few concepts which will be used in the sequel. In Section 3, we present a reduction of the logspace shortest path algorithm to biconnected graphs. In Section 4, we present a linear time logspace algorithm for parameter-constrained graphs. We conclude with some final remarks in Section 5.

## 2 Preliminaries

A *logspace algorithm* is an algorithm which uses  $O(\log n)$  working memory, where  $n$  is the size of the input. In addition, the input and output are respectively read-only and write-only, and do not count toward the space used. The *shortest path problem* requires finding a path between two given vertices  $s$  and  $t$  in a graph  $G$ , such that the sum of the weights of the edges constituting the path is as small as possible. In general, if  $s$  and  $t$  are not in the same connected component, or if the connected component containing  $s$  and  $t$  also contains a negative-weight cycle, the shortest path does not exist. For simplicity, we will assume there are no negative-weight cycles in  $G$ , although the proposed algorithms can be easily modified to detect (and terminate at) such cycles without any increase in overall complexity. We will also assume that  $G$  is encoded by its adjacency list, where vertices are labeled with the first  $n$  natural numbers. The  $j^{\text{th}}$  neighbor of vertex  $i$  is accessed with  $Adj(i, j)$  in  $O(1)$  time, and  $degree(i) = |Adj(i)|$ .

An *articulation point* in  $G$  is a vertex whose deletion increases the number of connected components of  $G$ . A *block* is a maximal subgraph of  $G$  which has no articulation points; if  $G$  has a single block, then  $G$  is *biconnected*. The *block tree*  $T$  of  $G$  is the bipartite graph with parts  $A$  and  $B$ , where  $A$  is the set of articulation points of  $G$  and  $B$  is the set of blocks of  $G$ ;  $a \in A$  is adjacent to  $b \in B$  if and only if  $b$  contains  $a$ . We define the *id* of a block in  $G$  to be (*largest, smallest*), where *largest* and *smallest* are the largest and smallest vertices in the block with respect to their labeling from 1 to  $n$ . Clearly, each block in  $G$  has a unique *id*. Note also that it is possible to lexicographically compare the *ids* of two or more blocks, i.e., if  $id_1 = (\ell_1, s_1)$  and  $id_2 = (\ell_2, s_2)$ , then  $id_1 > id_2$  if  $\ell_1 > \ell_2$  or if  $\ell_1 = \ell_2$  and  $s_1 > s_2$ .

Given numbers  $a_1$ ,  $a_2$ , and  $p$ , we define the *next* number after  $p$  as follows:

$$next(a_1, a_2, p) = \begin{cases} a_1 & \text{if } a_2 \leq p < a_1 \text{ or } p < a_1 \leq a_2 \text{ or } a_1 \leq a_2 \leq p \\ a_2 & \text{otherwise.} \end{cases}$$

We extend this definition to a list  $L$  of not necessarily distinct numbers by defining the *next* number in  $L$  after  $p$  to be a number in  $L$  larger than  $p$  by the smallest amount, or if no such number exists, to be the smallest number in  $L$ . The *next* number in  $L$  can be found with logspace and  $O(|L|)$  time, given sequential access to the elements of  $L$ , by repeatedly applying the *next* function.

### 3 Reducing the logspace shortest path problem to biconnected graphs

Let  $connected(H; v_1, v_2)$  be an implementation of Reingold's USTCON algorithm which takes in two vertices of a graph  $H$  and returns *true* if they belong to the same connected component, and *false* otherwise. Let  $pathInBlock(H; v_1, v_2)$  be a polynomial time, logspace oracle which takes in two vertices of a biconnected graph  $H$  and prints the shortest path between them.

Clearly, the encoding of a graph  $H$  can be reduced with logspace and polynomial time to the encoding of some induced subgraph  $H[S]$ . Thus, by transitivity and closure of reductions, the functions  $connected(H[S]; v_1, v_2)$  and  $pathInBlock(H[B]; v_1, v_2)$  can be used with logspace and polynomial time, where  $S$  and  $B$  are sets of vertices computed at runtime and  $H[B]$  is biconnected.

The  $connected$  function reduces the logspace shortest path problem to connected graphs. In this section, we will further reduce this problem to biconnected graphs, by presenting a logspace algorithm for finding the shortest path between two vertices in an arbitrary graph using the oracle  $pathInBlock$ .

#### 3.1 Constructing a logspace traversal function

Let  $G$  be a graph of order  $n$ , and  $v_1$  and  $v_2$  be two vertices that belong to the same block; the set of all vertices in this block will be referred to as  $block(v_1, v_2)$ . Using the  $connected$  function, is easy to construct a logspace function  $isInBlock(v_1, v_2, v)$  which returns *true* when  $v$  is part of  $block(v_1, v_2)$  and *false* otherwise; see Table 1 for pseudocode. This procedure can be used to access the vertices in  $block(v_1, v_2)$  sequentially. A similar procedure  $areInBlock(u, v)$  can be defined which returns *true* when  $u$  and  $v$  are in the same block, and *false* otherwise.

A vertex of  $G$  is an articulation point if and only if two of its neighbors are not in the same block. Thus, using the  $isInBlock$  function, we can construct a function  $isArticulation(v)$  which returns *true* when  $v$  is an articulation point and *false* otherwise; see Table 1 for pseudocode. We also define the function  $id(v_1, v_2)$ , which goes through the vertices of  $block(v_1, v_2)$  and returns (*largest*, *smallest*), where *largest* and *smallest* are respectively the largest and smallest vertices in  $block(v_1, v_2)$  according to their labeling.

Let  $p$  be an articulation point<sup>1</sup> in  $block(v_1, v_2)$ . To find the *next* articulation point in  $block(v_1, v_2)$  after  $p$ , we can create a function  $nextArticulation(v_1, v_2, p)$  which uses each articulation point in  $block(v_1, v_2)$  as a member of list  $L$  and applies the *next* function. Note that the vertices in  $L$  do not have to be stored, but can be generated one at a time; see Table 1 for pseudocode. Similarly, to identify the block containing  $p$  and having the *next id* after  $id(v_1, v_2)$ , we can create a function  $nextBlock(v_1, v_2, p)$  which uses the *ids* of the blocks identified by  $p$  and each of its neighbors as members of a list  $L$  and applies the *next*

<sup>1</sup> The subsequent definitions and functions remain valid when  $p$  is not an articulation point, and can be used in special cases, e.g., when  $G$  only has one block.

function. Note that the *ids* in  $L$  do not have to be stored but can be computed one at a time; see Table 1 for pseudocode.

Finally, given articulation point  $p$  and vertex  $v$  in the same block, we will call the component of  $G - \{block(v,p) \setminus \{p\}\}$  which contains  $p$  the *subgraph of  $G$  rooted at  $block(v,p)$  containing  $p$* , or *subgraph( $v,p$ )*. This subgraph can be traversed with logspace by starting from  $p$  and repeatedly moving to the *next* block and to the *next* articulation point until the starting block is reached again. This procedure indeed gives a traversal, since it corresponds to visiting the *next* neighbor in the block tree  $T$  of  $G$ , which generates an Euler subtour traversal (cf. [14]). In addition, during the traversal of *subgraph( $v,p$ )*, each vertex can be compared to a given vertex  $t$ , in order to determine whether the subgraph contains  $t$ . Thus, we can create a function *isInSubgraph( $v,p,t$ )* which returns *true* if  $t$  is in *subgraph( $v,p$ )* and *false* otherwise; see Table 1 for pseudocode.

---

**Table 1:** List of subroutines

---

```

function isInBlock( $v_1, v_2, v$ )
  if  $v = v_1$  or  $v = v_2$  then return true;
  for  $i = 1$  to  $n$ ,  $i \neq v_1$ ,  $i \neq v$  do
    if  $\neg connected(G - i; v_1, v)$  then return false;
  if  $\neg connected(G - v_1; v_2, v)$  then return false;
  return true;

function isArticulation( $v$ )
  for  $i = 2$  to degree( $v$ ) do
    if  $\neg isInBlock(v, Adj(v, i), Adj(v, i))$  then return true;
  return false;

function nextArticulation( $v_1, v_2, p$ )
   $a = p$ ;
  for  $v = 1$  to  $n$  do
    if isInBlock( $v_1, v_2, v$ ) and isArticulation( $v$ ) then  $a = next(a, v, p)$ ;
  return a;

function nextBlock( $v_1, v_2, p$ )
   $id^* = id(v_1, v_2)$ ;  $a = id(p, Adj(p, 1))$ ;  $i^* = 1$ ;
  for  $i = 2$  to degree( $p$ ) do
     $b = id(p, Adj(p, i))$ ;
    if  $a \leq id^* < b$  or  $id^* < b \leq a$  or  $b \leq a \leq id^*$  then  $a = b$ ;  $i^* = i$ ;
  return  $Adj(p, i^*)$ ;

function isInSubgraph( $v, p, t$ )
   $v_1 = p$ ;  $v_2 = v$ ;  $id^* = id(v_1, v_2)$ ;  $v_2 = nextBlock(v_1, v_2, p)$ ;
  while  $id(v_1, v_2) \neq id^*$  do
     $p = nextArticulation(v_1, v_2, p)$ ;  $v = nextBlock(v_1, v_2, p)$ ;
    if isInBlock( $v_1, v_2, t$ ) then return true;
     $v_1 = p$ ;  $v_2 = v$ ;
  return false;

```

---

### 3.2 Main Algorithm

Using the subroutines outlined in the previous section and the oracle *pathInBlock*, we propose the following logspace algorithm for finding the shortest path in a graph  $G$ . The main idea is to print the shortest path one block at a time by locating  $t$  in one of the subgraphs rooted at the current block.

---

**Algorithm 1:** Shortest path using oracle

---

```

function shortestPath( $G; s, t$ )
  if  $\neg$ connected( $G; s, t$ ) then return “Path does not exist”;
   $current = s$ ;
  start
  if areInBlock( $current, t$ ) then
    return pathInBlock( $G[block(current, t)]; current, t$ );
  for  $i = 1$  to degree( $current$ ) do
    for  $potential = 1$  to  $n$  do
      if isInBlock( $current, Adj(current, i), potential$ ) and
        isArticulation( $potential$ ) and
        isInSubgraph( $current, potential, t$ ) then
          pathInBlock( $G[block(current, potential)]; current, potential$ );
           $current = potential$ ;
          goto start;

```

---

**Theorem 1.** *Algorithm 1 finds the correct shortest path between vertices  $s$  and  $t$  in graph  $G$  with logspace and polynomial time, using a shortest path oracle for biconnected graphs.*

*Proof.* Let  $p_0 = s$  and  $p_{\ell+1} = t$ ; the shortest path between  $p_0$  and  $p_{\ell+1}$  is  $P = p_0P_0p_1P_1 \dots p_{\ell}P_{\ell}p_{\ell+1}$ , where  $p_1, \dots, p_{\ell}$  are articulation points and  $P_0, \dots, P_{\ell}$  are (possibly empty) subpaths which contain no articulation points. Let  $b_i = block(p_i, p_{i+1})$  for  $0 \leq i \leq \ell$ , so that  $pathInBlock(G[b_i]; p_i, p_{i+1}) = p_iP_i p_{i+1}$ .

Suppose the subpath  $p_0P_0 \dots p_i$ ,  $i \geq 0$ , has already been printed and that the vertex  $p_i$  is stored in memory. In each iteration of the main loop, the function  $isInSubgraph(p_i, p, t)$  returns *true* only for  $p = p_{i+1}$  when run for all articulation points  $p$  in all blocks containing  $p_i$ . The function  $pathInBlock(G[b_i], p_i, p_{i+1})$  is then used to print  $P_{i+1}$  and  $p_{i+1}$ . Finally,  $p_i$  is replaced in memory by  $p_{i+1}$ , and this procedure is repeated until  $p_{\ell+1}$  is reached. Since the main loop is entered only if the shortest path is of finite length, the algorithm terminates, and since each subpath printed is between two consecutive articulation points of  $P$ , the output of Algorithm 1 is the correct shortest path between  $s$  and  $t$ .

Since the *connected* function is logspace, the *isInBlock*, *isArticulation* and *isInSubgraph* functions are each logspace. Only a constant number of variables, each of size  $O(\log n)$ , are simultaneously stored in Algorithm 1, and every function call is to a logspace function (assuming the *pathInBlock* oracle is logspace); thus, the space complexity of Algorithm 1 is  $O(\log n)$ . Note that since the vertices in  $block(v_1, v_2)$  cannot be stored in memory simultaneously, a call to the

function  $pathInBlock(G[block(v_1, v_2)], v_1, v_2)$  needs to be realized by a logspace reduction, i.e., the vertices  $v_1$  and  $v_2$  are stored, and whenever the function  $pathInBlock$  needs to access an entry of the adjacency list of  $G[V(block(v_1, v_2))]$ , it recomputes it by going through the vertices of  $G$  and using the function  $isInBlock$ .

Similarly, since the *connected* function uses polynomial time, the *isInBlock*, *isArticulation* and *isInSubgraph* functions each use polynomial time. The main loop is executed at most  $O(n)$  times, and each iteration calls a constant number of polynomial time functions (assuming the *pathInBlock* oracle uses polynomial time); thus, the time complexity of Algorithm 1 is  $O(n^c)$  for some constant  $c$ .  $\square$

## 4 Linear time logspace algorithm for parametrically constrained graphs

Let  $BellmanFord(H; v_1, v_2)$  be an implementation of the Bellman-Ford shortest path algorithm [3] which takes in two vertices of a graph  $H$  and prints out the shortest path between them. Let  $HopcroftTarjan(H)$  be an implementation of Hopcroft and Tarjan's algorithm [8] which returns all blocks and articulation points of a graph  $H$ . If the size of  $H$  is bounded by a constant,  $BellmanFord$  and  $HopcroftTarjan$  can each be used with constant time and a constant number of memory cells.

Let  $G$  be a graph of order  $n$  with maximum vertex degree  $\Delta$  and maximum biconnected component size  $k$ . We will regard  $\Delta$  and  $k$  as fixed constants, independent of  $n$ . Using these constraints and some additional computational techniques, we will reformulate Algorithm 1 as a linear-time logspace shortest path algorithm which does not rely on an oracle. Asymptotically, both the time and space requirements of this algorithm are the best possible and cannot be improved; see Corollary 1 for more information.

### 4.1 Constructing a linear time logspace traversal function

By the assumption on the structure of  $G$ , the number of vertices at distance at most  $k$  from a specified vertex  $v$  is bounded by  $\lfloor \frac{\Delta^{k+1}-1}{\Delta-1} \rfloor$ . Thus, any operations on a subgraph induced by such a set of vertices can be performed with constant time and a constant number of memory cells, each with size  $O(\log n)$ ; note that since each vertex of  $G$  has a bounded number of neighbors,  $G[S]$  can be found in constant time for any set  $S$  of bounded size. In particular, we can construct a function  $blocksContaining(v)$  which uses  $HopcroftTarjan$  to return all blocks containing a given vertex  $v$  and all articulation points in these blocks; see below for pseudocode.

Using the set of blocks and articulation points given by the  $blocksContaining$  function, we can define functions  $isInBlock(v_1, v_2, v)$ ,  $areInBlock(u, v)$ ,  $isArticulation(v)$ ,  $id(v_1, v_2)$ ,  $nextArticulation(v_1, v_2, p)$ , and  $nextBlock(v_1, v_2, p)$  analogous to the ones described in Section 3, each of which uses  $O(\log n)$  space and  $O(1)$  time. We can also construct an analogue of  $isInSubgraph(v, p, t)$ , which

---

**Subroutine:** Finding all blocks containing  $v$ , and their articulation points

---

```

function blocksContaining( $v$ )
   $S = \{v\}$ ;
  for  $i = 1$  to  $k$  do  $S = \bigcup_{v \in S} \text{Adj}(v)$ ;
   $(\mathcal{B}, \mathcal{A}) = \text{HopcroftTarjan}(G[S])$ ;
  //  $\mathcal{B}$  is set of blocks,  $\mathcal{A}$  is set of articulation points of  $G[S]$ 
  return  $(B = \{b \in \mathcal{B} : b \cap \text{Adj}(v) \neq \emptyset\}, A = \mathcal{A} \cap \bigcup_{b \in B} b)$ ;

```

---

uses time proportional to the size of  $\text{subgraph}(v, p)$ ; in particular, the time for traversing the entire graph  $G$  via an Euler tour of its block tree is  $O(n)$  (provided  $G$  is connected) since there are  $O(n)$  calls to the *nextArticulation* function and  $O(n)$  calls to the *nextBlock* function.

Finally, it will be convenient to define the following functions: *adjacentPoints*( $v_1, v_2$ ) which returns the set of articulation points belonging to blocks containing  $v_2$  but not  $v_1$  if  $v_1 \neq v_2$  and the set of articulation points belonging to blocks containing  $v_2$  if  $v_1 = v_2$  (this function is slight modification of *blocksContaining*); *last*( $L$ ) which returns the last element of a list  $L$ ; *traverseComponent*( $s, t$ ) which traverses the component containing a vertex  $s$  and returns *true* if  $t$  is in the same component and *false* otherwise (this function is identical to *isInSubgraph*, with a slight modification in the stopping condition).

## 4.2 Linear time logspace shortest path algorithm

We now present a modified version of Algorithm 1, which uses the subroutines outlined in the previous section as well as some additional computational techniques such as “simulated parallelization” (introduced by Asano et al. [1]) aimed at reducing its runtime.

---

**Algorithm 2:** Shortest path in parameter constrained graph

---

```

function shortestPath( $G; s, t$ )
  if  $\neg \text{traverseComponent}(s, t)$  then return “Path does not exist”;
   $previous = s; current = s$ ;
  start
  if  $\text{areInBlock}(current, t)$  then
    return  $\text{BellmanFord}(G[\text{block}(current, t)]; current, t)$ ;
  for  $potential \in L := \text{adjacentPoints}(previous, current)$  using two
  pointers in serial do
    if  $potential = \text{last}(L)$  or  $\text{isInSubgraph}(current, potential, t)$  then
       $\text{BellmanFord}(G[\text{block}(current, potential)]; current, potential)$ ;
       $previous = current; current = potential$ ;
      goto start;

```

---

**Theorem 2.** *Algorithm 2 finds the correct shortest path between vertices  $s$  and  $t$  in graph  $G$  with bounded degree and biconnected component size using logspace and linear time.*

*Proof.* Using the notation in the proof of Theorem 1, suppose the subpath  $p_0P_0\dots p_i$  has already been printed;  $p_{\ell+1}$  cannot be in  $subgraph(p_i, p_{i-1})$ , so there is no need to run  $isInSubgraph(p_i, p_{i-1}, t)$ . Thus,  $adjacentPoints(p_{i-1}, p_i)$  is the set of feasible articulation points. Moreover, if  $p_{\ell+1}$  is not in  $subgraph(p_i, p)$  for all-but-one feasible articulation points, then the last of these must be  $p_{i+1}$  and there is no need to run  $isInSubgraph(p_i, p_{i+1}, t)$ . Finally, two subgraphs rooted at  $block(p_{i-1}, p_i)$  can be traversed concurrently with the technique of simulated parallelization: instead of traversing the feasible subgraphs one-after-another, we maintain two copies of the  $isInSubgraph$  function and use them to simultaneously traverse two subgraphs. We do this in serial (without the use of a parallel processor) by iteratively advancing each copy of the function in turn; if one subgraph is traversed, the corresponding copy of the function terminates and another copy is initiated to traverse the next unexplored subgraph. Thus, Algorithm 2 is structurally identical to Algorithm 1<sup>2</sup> and prints the correct shortest path between  $s$  and  $t$ .

Only a constant number of variables, each of size  $O(\log n)$ , are simultaneously used in Algorithm 2, and every function call is to a logspace function; moreover, keeping track of the internal states of two logspace functions can be done with logspace, so the space complexity of Algorithm 2 is  $O(\log n)$ .

Finally, to verify the time complexity, note that by traversing two subgraphs at once, we can deduce which subgraph contains  $t$  in the time it takes to traverse all subgraphs which do *not* contain  $t$  or  $s$ . Thus, each subgraph rooted at  $block(p_i, p_{i+1})$ ,  $0 \leq i \leq \ell$ , which does not contain  $t$  or  $s$  will be traversed at most once, so the time needed to print the shortest path is of the same order as the time needed to traverse  $G$  once.  $\square$

**Corollary 1.** *The time and space complexity of Algorithm 2 is the best possible for the class of graphs considered.*

*Proof.* Let  $G$  be a graph of order  $n$ ; the shortest path between two vertices in  $G$  may be of length  $\Omega(n)$  so any shortest path algorithm will require at least  $\Omega(n)$  time to print the path. Moreover, a pointer to an entry in the adjacency list of  $G$  has size  $\Omega(\log n)$ , so printing each edge of the shortest path requires at least  $\Omega(n)$  space.  $\square$

## 5 Conclusion

We have reduced the logspace shortest path problem to biconnected graphs using techniques such as computing instead of storing, transitivity of logspace reductions, and Reingold’s USTCON result. We have also proposed a linear time

<sup>2</sup> Indeed each of the described modifications can be implemented in Algorithm 1 as well, but would not make a significant difference in its time complexity.



logspace shortest path algorithm for graphs with bounded degree and biconnected component size, using techniques such as simulated parallelization and constant-time and -space calls to functions over graphs with bounded size.

Future work will be aimed at further reducing the logspace shortest path problem to triconnected graphs using SPQR-tree decomposition, and to  $k$ -connected graphs using branch decomposition or the decomposition of Holberg [7]. Another direction for future work will be to generalize Algorithm 2 by removing or relaxing the restrictions on vertex degree and biconnected component size.

## Acknowledgements

This material is based upon work supported by the National Science Foundation under Grants No. 1450681, CMMI-1300477, and CMMI-1404864.

## References

1. Asano, T., Mulzer, W., Wang, Y., Constant work-space algorithms for shortest paths in trees and simple polygons. *J. of Graph Algorithms and Applications* 15(5), 569–586 (2011)
2. Barnes, G., Feige, U., Short random walks on graphs. *Proc. 25th Annual ACM Symposium of the Theory of Computing (STOC)*, 728–737 (1993)
3. Bellman, R., On a routing problem. *Quarterly of Applied Mathematics* 16 87–90 (1958)
4. Brimkov, B., Hicks, I.V., Memory efficient algorithms for cactus graphs and block graphs. *Discrete Applied Math.* In Press, doi:10.1016/j.dam.2015.10.032 (2015)
5. Das, B., Datta, S., Nimbhorkar, P., Log-space algorithms for paths and matchings in  $k$ -trees. *Theory of Computing Systems* 53 (4) 669–689 (2013)
6. Feige, U., A randomized time-space trade-off of  $\tilde{O}(mR)$  for USTCON. *Proc. 34th Annual Symposium on Foundations of Computer Science (FOCS)*, 238–246 (1993)
7. Holberg, W., The decomposition of graphs into  $k$ -connected components. *Discrete Mathematics* 109(1), 133–145 (1992)
8. Hopcroft, J. and Tarjan, R., Algorithm 447: efficient algorithms for graph manipulation. *Communications of the ACM* 16(6), 372–378 (1973)
9. Jakoby, A., Tantau, T., Logspace algorithms for computing shortest and longest paths in series-parallel graphs. *FSTTCS 2007: Foundations of Software Technology and Theoretical Computer Science*. Springer Berlin Heidelberg, 216–227 (2007)
10. Jakoby, A., Liskiewicz, M., Reischuk, R., Space efficient algorithms for series-parallel graphs. *J. of Algorithms* 60, 85–114 (2006)
11. Kosowski, A., Faster walks in graphs: a  $O(n^2)$  time-space trade-off for undirected st connectivity. *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM (2013)
12. Munro, I.J., Ramirez, R. J., Technical note—reducing space requirements for shortest path problems. *Operations Research* 30.5, 1009–1013 (1982)
13. Reingold, O., Undirected connectivity in log-space. *J. ACM* 55(4), Art. 17 (2008)
14. Tarjan, R.E. and Vishkin, U., Finding biconnected components and computing tree functions in logarithmic parallel time. *Proceedings of FOCS*, 12–20 (1984)