

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Frameworks for designing in-place graph algorithms [☆]

Sankardeep Chakraborty ^a, Anish Mukherjee ^b, Venkatesh Raman ^c,
Srinivasa Rao Satti ^{d,*}

^a University of Tokyo, Japan^b University of Warsaw, Poland^c The Institute of Mathematical Sciences, HBNI, India^d Norwegian University of Science and Technology, Norway

ARTICLE INFO

Article history:

Received 9 February 2020

Received in revised form 23 June 2021

Accepted 18 July 2021

Available online 28 July 2021

Keywords:

Graph algorithms

In-place algorithms

Implicit model

ROM model

BFS

DFS

Topological sorting

ABSTRACT

Read-only memory (ROM) model is a classical model of computation to study time-space tradeoffs of algorithms. More recently, several graph algorithms have been studied under ROM model. In this paper, we study graph algorithms under two different relaxations of ROM model, referred to as the implicit and rotate models, and show that these simple relaxations allow us to implement fundamental graph search methods like BFS and DFS more space efficiently than in ROM model. All our algorithms are simple but quite subtle, and we believe that these models are practical enough to spur interest for other graph problems in these models.

© 2021 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Motivated by the rapid growth of huge datasets (“big data”), space efficient algorithms are becoming more important than ever before. The proliferation of specialized handheld devices and embedded systems that have a limited supply of memory provide another motivation to consider space efficient algorithms. To design *space-efficient* algorithms in general, several models of computation have been proposed. Among them, the following two computational models have received considerable attention in the literature.

- In the *read-only* memory (ROM) model, we assume that the input is given in a read-only memory. The output of an algorithm is written on to a separate write-only memory, and the output can not be read or modified again. In addition to the input and output media, a limited random access workspace is available. Early work on this model was on designing lower bounds [17,19,20], for designing algorithms for selection and sorting [35,53,61,72,73,75] and problems

[☆] These results were announced in preliminary form in the proceedings of 26th Annual European Symposium on Algorithms (ESA 2018) LIPIcs, Volume 112, pages 13:1–13:16 [28]. The work of the first author was supported by MEXT Quantum Leap Flagship Program (MEXT Q-LEAP) Grant Number JPMXS0120319794. The work of the second author was supported in part by ERC CoG project TUgBOAT number 772346.

* Corresponding author.

E-mail addresses: sankardeep.chakraborty@gmail.com (S. Chakraborty), anish343@gmail.com (A. Mukherjee), vraman@imsc.res.in (V. Raman), srinivasa.r.satti@ntnu.no (S.R. Satti).

<https://doi.org/10.1016/j.jcss.2021.07.004>

0022-0000/© 2021 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

in computational geometry [8,11,15,34,40]. Recently there has been interest on space-efficient graph algorithms [9,14,13,25,30,33,32,31,27,29,52,62,63].

- In the *in-place* model, the input elements are given in an array, and the algorithm may use the input array as working space. Hence the algorithm may modify the array during its execution. After the execution, all the input elements should be present in the array (maybe in a permuted order) and the output maybe put in the same array or sent to an output stream. The amount of extra space usage during the entire execution of the algorithm is limited to $O(\lg n)$ bits.¹ A prominent example of an in-place algorithm is the classic heap-sort. Other than in-place sorting [60], searching [58,70] and selection [67] algorithms, many in-place algorithms have been designed in areas such as computational geometry [21] and stringology [59].

Apart from these models, researchers have also considered (semi)-streaming models [5,56,72] for designing space-efficient algorithms. Very recently the following two new models were introduced in the literature with the same objective.

- Chan et al. [36] introduced the *restore* model which is a more relaxed version of read-only memory (and a restricted version of the in-place model), where the input is allowed to be modified, but at the end of the computation, the input has to be restored to its original form. They also gave space efficient algorithms for selection and sorting on integer arrays in this model. This has motivation, for example, in scenarios where the input (in its original form) is required by some other application.
- Buhrman et al. [22,23,66] introduced and studied the *catalytic-space* model where a small amount (typically $O(\lg n)$ bits) of clean space is provided along with some large additional auxiliary space, with the condition that the additional space is initially in an arbitrary, possibly incompressible, state and must be returned to this state when the computation is finished. The input is assumed to be given in ROM. Thus this model can be thought of as having an auxiliary storage that needs to be 'restored' in contrast to the model by Chan et al. [36] where the input array has to be 'restored'. They show various interesting complexity theoretic consequences in this model and designed significantly better (in terms of space) algorithms in comparison with the ROM model for a few combinatorial problems.

1.1. Previous work in space efficient graph algorithms

Even though these models were introduced in the literature with the aim of designing and/or implementing various algorithms space efficiently, *space efficient graph algorithms* have been designed only in the (semi)-streaming and the ROM model. In the streaming and semi-streaming models, researchers have studied several basic and fundamental algorithmic problems such as connectivity, minimum spanning tree, matching. See [69] for a comprehensive survey in this field. Research on these two models (i.e., streaming and semi-streaming) is relatively new and has been going on for last decade or so whereas the study in ROM model could be traced back to almost 40 years. In fact there is already a rich history of designing space efficient algorithms in the read-only memory model. The complexity class L (also known as DLOGSPACE) is the class containing decision problems that can be solved by a deterministic Turing machine using only logarithmic amount of work space for computation. There are several important algorithmic results [43,48–50] for this class, the most celebrated being Reingold's method [78] for checking st -reachability in an undirected graph, i.e., to determine if there is a path between two given vertices s and t . NL is the non-deterministic analogue of L and it is known that the st -reachability problem for *directed* graphs is NL-complete (with respect to log-space reductions). Using Savitch's algorithm [7], this problem can be solved in $n^{O(\lg n)}$ time using $O(\lg^2 n)$ bits of extra space. Savitch's algorithm is very space efficient but its running time is superpolynomial.

Among the deterministic algorithms running in polynomial time for directed st -reachability, the most space efficient algorithm is due to Barnes et al. [16] who gave a slightly sublinear space (using $n/2^{\Theta(\sqrt{\lg n})}$ bits) algorithm for this problem running in polynomial time. We know of no better polynomial time algorithm for this problem with better space bound. Moreover, the space used by this algorithm matches a lower bound on space for solving directed st -reachability on a restricted model of computation called Node Naming Jumping Automata on Graphs (NNJAG's) [37,47]. This model was introduced especially for the study of directed st -reachability and most of the known sublinear space algorithms for this problem can be implemented on it. Thus, to design any polynomial time ROM algorithm taking space less than $n/2^{\Theta(\sqrt{\lg n})}$ bits for st -reachability requires radically new ideas. Recently there has been some improvement in the space bound for some special classes of graphs like planar and H-minor free graphs [10,24]. A drawback for all these algorithms using small space i.e., sublinear number of bits, is that their running time is some polynomial of high degree. This is not surprising as Tompa [82] showed that for directed st -reachability, if the number of bits available is $o(n)$ then some natural algorithmic approaches to the problem require super-polynomial time.

Motivated by these impossibility results from complexity theory and inspired by the practical applications of these fundamental graph algorithms, recently there has been a surge of interest in improving the space complexity of the fundamental graph algorithms without paying too much penalty in the running time i.e., reducing the working space of the classical

¹ We use \lg to denote logarithm to the base 2.

graph algorithms to $O(n)$ bits with little or no penalty in running time. Thus the goal is to design space-efficient yet reasonably time-efficient graph algorithms on the ROM model. Generally most of the classical linear time graph algorithms take $O(n)$ words or equivalently $O(n \lg n)$ bits of space. Recently Asano et al. [9] gave an $O(m \lg n)$ time algorithm using $O(n)$ bits, and another implementation taking $n + o(n)$ bits, using Reingold's or Barnes et al.'s reachability algorithm, and hence have high polynomial running time. Later, time bound was improved to $O(m \lg \lg n)$ still using $O(n)$ bits in [52]. For sparse graphs, the time bound is further improved in [14,30] to optimal $O(m)$ using still $O(n)$ bits of space. Improving on the classical linear time implementation of BFS which uses $O(n \lg n)$ bits of space, recent space efficient algorithms [14,52,62] have resulted in a linear time algorithm using $n \lg 3 + o(n)$ bits. We know of no algorithm for BFS using $n + o(n)$ bits and $O(m \lg^c n)$ (or even $O(mn)$) time for some constant c in ROM. The only BFS algorithm taking sublinear space uses $n/2^{\Theta(\sqrt{\lg n})}$ bits [16] and has a high polynomial runtime. A few other space efficient algorithms for fundamental graph problems like checking strong connectivity [52], biconnectivity and performing st -numbering [30], recognizing chordal and outerplanar graphs [33,63] were also designed very recently.

1.2. Organization of the paper

In what follows, we first define both the models and summarize our main results. We also briefly discuss our main techniques and the consequences of our results in the remaining part of this section. Section 2 describes DFS algorithms for the rotate model, and this is followed by the BFS algorithms for the rotate model in Section 3. Section 4 describes an efficient method to simulate any the rotate model algorithm in the implicit model. In Section 5, we provide the DFS algorithm for the rotate model, and this is followed the implicit model DFS algorithm in Section 6. Section 7 discusses the algorithms for minimum spanning tree in both the rotate model and the implicit model. In Section 8, we provide various consequences and implications of our results. Finally, Section 9 concludes with some possible future directions.

1.3. In-place model for graph algorithms

In order to break these inherent space bound barriers and obtain reasonably time and space efficient graph algorithms, we want to relax the limitations of ROM. The most natural and obvious candidate in this regard is the classical *in-place* model. Thus our main objective is to initiate a systematic study of efficient in-place (i.e., using $O(\lg n)$ bits of extra space) algorithms for graph problems. To the best of our knowledge, this has not been done in the literature before. Our first goal towards this is to properly define models for the in-place graph algorithms. As in the case of standard in-place model, we need to ensure that the graph (adjacency) structure remains intact throughout the algorithm. Let $G = (V, E)$ be the input graph with $n = |V|$, $m = |E|$, and assume that the vertex set V of G is the set $V = \{1, 2, \dots, n\}$. To describe these models, we assume that the input graph representation consists of two parts: (i) an array V of length n , where $V[i]$ stores a pointer to the adjacency list of vertex i , and (ii) a list of singly linked lists, where the i -th list consists of a singly linked list containing all the neighbors of vertex i with $V[i]$ pointing to the head of the list. In the ROM model, we assume that both these components cannot be modified. In our relaxed models, we assume that one of these components can be modified in a limited way. This gives rise to two different models which we define next.

Implicit model: The most natural analogue of in-place model allows any two elements in the adjacency list of a vertex to be swapped (in constant time assuming that we have access to the nodes storing those elements in the singly linked list). The adjacency "structure" of the representation does not change; only the values stored can be swapped. (One may restrict this further to allow only elements in adjacent nodes to be swapped. Most of our algorithms work with this restriction.) We call it the implicit model inspired by the notion of *implicit data structures* [70]. We introduce and develop algorithms for another relaxed model which we call the rotate model.

Rotate model: In this model, we assume that only the pointers stored in the array V can be modified, that too in a limited way - to point to any node in the adjacency list, instead of always pointing to the first node. In space-efficient setting, since we do not have additional space to store a pointer to the beginning of the adjacency list explicitly, we assume that the second component of the graph representation consists of a list of circular linked lists (instead of singly linked lists) - i.e., the last node in the adjacency list of each vertex points to the first node (instead of storing a null pointer). See the Fig. 1 to get a better visual description. We call the element pointed to by the pointer as the front of the list, and a unit cost rotate operation changes the element pointed to by the pointer to the next element in the list.

Thus the rotate model corresponds to keeping the adjacency lists in read-only memory and allowing (limited) updates on the pointer array that points to these lists. And, the implicit model corresponds to the reverse case, where we keep the pointer array in read-only memory and allow swaps on the adjacency lists/arrays. A third alternative especially for the implicit model is to assume that the input graph is represented as an adjacency array, i.e., adjacency lists are stored as arrays instead of singly linked lists (see [30,52,63] for some results using this model); and we allow here that any two elements in the adjacency array can be swapped. In this model, some of our algorithms have improved performance in time.

Tarjan [80] considered 'space efficient implementations of graph search methods', and proposed several implementations for BFS and DFS using $2n + m$ bits of extra space. The input representation is modified during the execution of these algorithms. Thus his model is analogous to the implicit model described above, although the space usage of these algorithms




Fig. 1. (a) An undirected graph G with 5 vertices and 8 edges. (b) A circular list representation of G . To avoid cluttering the picture, we draw the vertices and the pointers to the next node separately as opposed to a single node having two different fields in the circular list. (c) An illustration of a single clockwise rotation in the circular list of vertex 4.

is significantly high. Also, his model allows the graph structure to be ‘modified’ during the execution of the algorithm, although it is restored at the end of the algorithm.

1.4. Definitions, computational complexity and notations

We study some basic and fundamental graph problems in these models. In what follows we provide the definitions and state the computational complexity of some of these problems. For the DFS problem, there have been two versions studied in the literature. In the *lexicographically smallest DFS* or *lex-DFS* problem, when DFS looks for an unvisited vertex to visit in an adjacency list, it picks the “first” unvisited vertex where the “first” is with respect to the appearance order in the adjacency list. The resulting DFS tree will be unique. The decision version of this problem can be defined as given vertices u and v , is u visited before v in the lex-DFS order? In contrast to lex-DFS, an algorithm that outputs *some* DFS numbering of a given graph, treats an adjacency list as a set, ignoring the order of appearance of vertices in it, and outputs a vertex ordering T such that there exists *some* adjacency ordering R such that T is the DFS numbering with respect to R . We say that such a DFS algorithm performs *general-DFS*. Reif [77] has shown that lex-DFS is P-complete (with respect to log-space reductions) implying that a logspace algorithm for lex-DFS results in the collapse of complexity classes P and L. Anderson et al. [6] have shown that even computing the leftmost root-to-leaf path of the lex-DFS tree is P-complete. For many years, these results seemed to imply that the general-DFS problem, that is, the computation of any DFS tree is also inherently sequential. However, Aggarwal et al. [3,4] proved that the general-DFS problem can be solved much more efficiently, and it is in RNC. Whether the general-DFS problem is in NC is still open.

As is standard in the design of space-efficient algorithms [14,52], while working with directed graphs, we assume that the graphs are given as in/out (circular) adjacency lists, i.e., for a vertex v , we have the (circular) lists of both in-neighbors and out-neighbors of v . We assume the word RAM model of computation where the machine consists of words of size w in $\Omega(\lg n)$ bits and any logical, arithmetic and bitwise operation involving a constant number of words takes $O(1)$ time. We count space in terms of number of *extra* bits used by the algorithm other than the input, and this quantity is referred as “extra space” and “space” interchangeably throughout the paper. By a path of length d , we mean a simple path on d edges. By $deg(x)$ we mean the degree of the vertex x . In directed graphs, it should be clear from the context whether that denotes out-degree or in-degree. By a BFS/DFS traversal of the input graph G , as in [9,14,30,52,62], we refer to reporting the vertices of G in the BFS/DFS ordering, i.e., in the order in which the vertices are visited for the first time.

1.5. Our results

Rotate Model: For DFS, in the rotate model, we show the following in Sections 2.1, 2.3 and 2.5.

Theorem 1. Let G be a directed or an undirected graph, and $\ell \leq n$ be the maximum depth of the DFS tree starting at a source vertex s . Then in the rotate model, the vertices of G can be output in

- (a) the lex-DFS order in $O(m+n)$ time using $n \lg 3 + O(\lg^2 n)$ bits,
- (b) a general-DFS order in $O(m+n)$ time using $n + O(\lg n)$ bits, and
- (c) a general-DFS order in $O(m^2/n + m\ell)$ time for an undirected graph and in $O(m(n + \ell^2))$ time for directed graphs using $O(\lg n)$ bits. For this algorithm, we assume that s can reach all other vertices.

This is followed by the BFS algorithms where, in the rotate model, we show the following in Sections 3.1 and 3.2.

Theorem 2. Let G be a directed or an undirected graph, and ℓ be the depth of the BFS tree starting at the source vertex s . Then in the rotate model, the vertices of G can be output in a BFS order in

- (a) $O(m + n\ell^2)$ time using $n + O(\lg n)$ bits, and
- (b) $O(m\ell + n\ell^2)$ time using $O(\lg n)$ bits. Here we assume that the source vertex s can reach all other vertices.

Implicit Model: In the implicit model, we obtain polynomial time implementations for lex-DFS and general-DFS using $O(\lg n)$ bits. For lex-DFS, this is conjectured to be unlikely in ROM as the decision problem is P-complete [77]. In particular, we show the following in Section 5.

Theorem 3. Let G be a directed or an undirected graph with a source vertex s and $\ell \leq n$ be the maximum depth of the DFS tree starting at s that can reach all other vertices. Then in the implicit model, using $O(\lg n)$ bits the vertices of G can be output in

- (a) the lex-DFS order in $O(m^3/n^2 + \ell m^2/n)$ time if G is given in adjacency list and in $O(m^2 \lg n/n)$ time if G is given in adjacency array for undirected graphs. For directed graphs our algorithm takes $O(m^2(n + \ell^2)/n)$ time if G is given in adjacency list and $O(m \lg n(n + \ell^2))$ time if G is given in adjacency array;
- (b) a general-DFS traversal order in $O(m^2/n)$ time if the input graph G is given in an adjacency list and in $O(m^2(\lg n)/n + m\ell \lg n)$ time if it is given in an adjacency array.

In the implicit model, we can match the runtime of BFS from rotate model, and do better in some special classes of graphs. In particular, we show the following in Section 6.

Theorem 4. Let G be a directed or an undirected graph with a source vertex that can reach all other vertices by a distance of at most ℓ . Then in the implicit model, using $O(\lg n)$ bits the vertices of G can be output in a BFS order in

- (a) $O(m + n\ell^2)$ time;
- (b) the runtime can be improved to $O(m + n\ell)$ time if there are no degree 2 vertices;
- (c) the runtime can be improved to $O(m)$ if the degree of every vertex is at least $2 \lg n + 3$.

In sharp contrast, for space efficient algorithms for DFS in ROM, the landscape looks markedly different. To the best of our knowledge, there are no DFS algorithms in general graphs in ROM that use $O(n^{1-\epsilon})$ bits. In fact, an implementation of DFS taking cn bits for $c < 1$ has been proposed as an open problem by Asano et al. [9]. Similar to DFS, to the best of our knowledge, there are no polynomial time BFS algorithms in ROM that use even $O(n^{1-\epsilon})$ bits. On the other hand, we don't hope to have a BFS algorithm (for both undirected and directed graphs) using $O(\lg n)$ bits in ROM as the shortest path distance problem which is NL-complete [7] reduces to this.

Minimum Spanning Tree (MST). Moving on from DFS and BFS, we also study the problem of reporting a minimum spanning tree (MST) of a given undirected connected graph G . We show the following result in Section 7.

Theorem 5. A minimum spanning forest of a given undirected weighted graph G can be found using $O(\lg n)$ bits and in

- (a) $O(mn)$ time in the rotate model,
- (b) $O(mn^2)$ time in the implicit model if G is given in an adjacency list, and
- (b) $O(mn \lg n)$ time in the implicit model when G is represented in an adjacency array.

Note that by the results of [76,78], we already know logspace algorithms for MST in ROM but again the drawback of those algorithms is their large time complexity. On the other hand, our algorithms are simple, and have relatively small polynomial running time, making them an appealing choice in applications with strict space constraints.

1.6. Techniques

Our implementations follow (variations of) the classical algorithms for BFS and DFS that use three colors, but avoid the use of stack (for DFS) and queue (for BFS). In the rotate model, we first observe that in the usual search algorithms one can dispense with the extra data structure space of pointers maintaining the search tree (while retaining the linear number of bits and a single bit per vertex in place of the full unvisited/visited/explored array) simply by rotating each circular adjacency lists to move the parent or a (typically the currently explored) child to the beginning of the list to help navigate through the tree during the forward or the backtracking step, i.e. by changing the pointer from the vertex to the list of its adjacencies by one node at a time. This retains the basic efficiency of the search strategies. The nice part of this strategy is that the total number of rotations also can be bounded. To reduce the extra space from linear to logarithmic, it is noted that one can follow the vertices based on the current rotations at each vertex to determine the visited status of a vertex, i.e. these algorithms use the rotate operation in a non-trivial way to move elements within the lists to determine the color of the vertices as well. However, the drawback is that to do so could require moving up (or down) the full height of the implicit search tree. This yields super-linear rather than (near-) linear time algorithms.

In the implicit model, we use the classical *bit encoding* trick used in the development of the implicit data structures [70]. We encode one (or two) bit(s) using a sequence of two (or three respectively) distinct numbers. To encode a single bit b using two distinct values x and y with $x < y$, we store the sequence x, y if $b = 0$, and y, x otherwise. Similarly, permuting three distinct values x, y, z with $x < y < z$, we can represent six combinations. We can choose any of the four combinations to represent up to 4 colors (i.e. two bits). Generalizing this further, we can encode a pointer taking $\lg n$ bits using $2 \lg n$ distinct elements where reading or updating a bit takes constant time, and reading or updating a pointer takes $O(\lg n)$ time. This also is the reason for the requirement of vertices with (high) degree at least 3 or $2 \lg n + 3$ for faster algorithms, which will become clear in the description of the algorithms.

1.7. Consequences of our BFS and DFS results

There are many interesting and surprising consequences of our results for BFS and DFS in both the rotate and implicit model. In what follows, we mention a few of them.

- For *directed st-reachability*, as mentioned previously, the most space efficient polynomial time algorithm [16] uses $n/2^{\Theta(\sqrt{\lg n})}$ bits. In sharp contrast, we obtain efficient (timewise) log-space algorithms for this problem in both the rotate and implicit models (as a corollary of our directed graph DFS/BFS results). In terms of workspace this is exponentially better than the best known polynomial time algorithm [16] for this problem in ROM. For us, this provides one of the main motivations to study this model. A somewhat incomparable result is obtained recently by Buhrman et al. [22,66] where they designed an algorithm for *directed st-reachability* on catalytic Turing machines in space $O(\lg n)$ with catalytic space $O(n^2 \lg n)$ and time $O(n^9)$.
- Problems like *directed st-reachability* [7] and *distance* [79] (which asks whether a given graph contains a path of length at most k from s to t) are NL-complete, i.e., no deterministic log-space algorithm is known for these problems. But in our (both the rotate and implicit) models, we design log-space algorithms for them. Assuming $L \neq NL$, these results show that probably both our models with log-space are stronger than L.
- Since the decision version of the lex-DFS problem (both in undirected and directed graphs) is P-complete [77], poly-logarithmic space algorithms are unlikely to exist in the ROM model for the lex-DFS problem. But we show an $O(\lg n)$ space algorithm in the implicit model for lex-DFS. This implies that, probably the implicit model is even more powerful than the rotate model. It could even be possible that every problem in P can be computed using log-space in the implicit model. A result of somewhat similar flavor is obtained recently Buhrman et al. [22,66] where they showed that any function in TC^1 can be computed using catalytic log-space, i.e., $TC^1 \subseteq CSPACE(\lg n)$. Note that TC^1 contains L, NL and even other classes that are conjectured to be different from L.
- Our bounds for BFS and DFS in the rotate and implicit models immediately imply (with some care) similar bounds, that are improvement over the best space bounds known so far in ROM, for many applications of DFS/BFS. Moreover, as described before, any algorithm in the rotate model can be implemented in ROM using extra $O(n \lg(m/n))$ bits. Thus, our linear time DFS algorithm in the rotate model can be implemented in ROM using $O(n \lg(m/n))$ bits, matching the bound of [30] for DFS. Using this DFS implementation, we can obtain improved space-efficient algorithms for various applications of DFS in ROM. This is discussed in Section 8.1.
- In Section 8.2 we present Logspace Approximation Scheme or LSAS ($(1 \pm \epsilon)$ approximation algorithm running in logspace for every $\epsilon > 0$) for a class of MSO-definable optimization problems which are amenable to Baker's method [12] in locally bounded treewidth graphs in both of our rotate and implicit models. No such algorithms are known in ROM as Baker's method requires to compute *distance* which is NL-complete. As BFS admits logspace algorithms in our models, we can design such LSAS for these problems here.
- For a large number of NP-hard graph problems, the best algorithms in ROM run in exponential time and polynomial space. We show that using just logarithmic amount of space, albeit using exponential time, we can design algorithms for those NP-hard problems in both of our models under some restrictions. This gives an exponential improvement over

the ROM space bounds for these problems. This is described in Section 8.3. In contrast, note that, no NP-hard problem can be solved in the ROM model using $O(\lg n)$ bits unless $P=NP$.

1.8. Recent developments

Shortly after our paper appeared in arXiv, Kammer and Sajenko [64] formulated another in-place model, and provided linear time algorithms for both DFS and BFS in their model. Note that our algorithms have super-linear running time. The efficiency of their algorithms, however, comes from stronger assumptions. More specifically, in our algorithms we did not use any bit-tricks. In other words, we assumed the input to be indivisible whereas Kammer and Sajenko do not have this restriction, and this allows them to do certain operations faster. Another important difference is that they assume the individual adjacency lists to be in sorted order, and they use this property crucially to pack more information by exploiting the redundancy. Again we do not have such input assumption and it would be really interesting to see if it is possible to design linear time algorithms for BFS and DFS in our models. We leave these as interesting open problems for future exploration.

2. DFS algorithms in the rotate model

In this section, we describe our space-efficient algorithms for DFS in the rotate model proving Theorem 1.

2.1. Lex-DFS using $n \lg 3 + O(\lg^2 n)$ bits—proof of Theorem 1(a) for undirected graphs

We begin by describing our algorithm for undirected graphs, and later mention the changes required for directed graphs. In the normal exploration of DFS (see for example, Cormen et al. [38]) we use three colors. Every vertex v is white initially while it has not been discovered yet, becomes gray when DFS discovers v for the first time, and is colored black when it is finished i.e., all its neighbors have been explored completely.

We maintain a color array C of length n that stores the color of each vertex at any point in the algorithm. In the rest of the paper, when we say we scan the adjacency list of some vertex v , what we mean is, we create a temporary pointer pointing to the current first element of the list and move this temporary pointer until we find the desired element. Once we get that element we actually rotate the list so that the desired element now is at the front of the list. We start DFS at the starting vertex, say s , changing its color from white to gray in the color array C . Then we scan the adjacency list of s to find the first white neighbor, say w . We keep rotating the list to bring w to the front of s 's adjacency list (as the one pointed to by the head $V[s]$), color w gray in the color array C and proceed to the next step (i.e. to explore w 's adjacency list). This is the first *forward* step of the algorithm. In general, at any step during the execution of the algorithm, whenever we arrive at a gray vertex u (including the case when u 's color is changed from white to gray in the current step), we scan u 's adjacency list to find the first white vertex. (i) If we find such a vertex, say v , then we rotate u 's list to make v as the first element, and change the color of v to gray. (ii) If we do not find any white vertex, then we change the color of u to black, and *backtrack* to its parent. To identify u 's parent, we use the following lemma.

Lemma 1. *Suppose w is a node that just became black. Then its parent p is the unique vertex in w 's adjacency list which is (a) gray and (b) whose current adjacency list has w in the first position.*

Proof. Among all the neighbors of w , some vertices are w 's children in the DFS tree, and the rest of them are w 's ancestors, and among the ancestors, exactly one vertex is w 's parent in the DFS tree. All the ancestors should have their currently explored (gray) child at the first position in their adjacency list; and this current child would be different from w for all the ancestors except p (as w was discovered from p). So, the second condition is violated for them. All of w 's children have been fully processed earlier and have been colored black, and hence the first condition is violated for them. Observe that, if w has a child, say k , which is a leaf in the DFS tree, it might happen that k also has w at the first position in its current adjacency list, but, fortunately, k is black while scanning w 's list. So for such vertices, the first condition gets violated. Only for w 's parent, which is p here, both the conditions are satisfied. \square

So, the parent can be found by scanning the w 's list, to find a neighbor p that is colored gray such that the first element in p 's list is w . This completes the description of the backtracking step. Once we backtrack to p , we find the next white vertex (as in the forward step) and continue until all the vertices of G are explored. Other than some constant number of variables, clearly the space usage is only for storing the color array C . Since C is of length n where each element has 3 possible values, C can be encoded using $n \lg 3 + O(\lg^2 n)$ bits, so that the i -th element in C can be read and updated in $O(1)$ time [46]. So overall space required is $n \lg 3 + O(\lg^2 n)$ bits. As the algorithm systematically brings a white vertex to the front, makes it gray, and moves it to the end after it becomes black, at most two full rotations of each of the list may happen (the second one to determine that there are no more white vertices) resulting in a linear time lex-DFS algorithm. We discuss the lex-DFS algorithm for the directed graphs below.

2.2. Lex-DFS using $n \lg 3 + O(\lg^2 n)$ bits–proof of Theorem 1(a) for directed graphs

Recall that we have access to both the in-adjacency and the out-adjacency lists for each vertex w in a directed graph G ; hence we can use these lists separately for performing two steps of DFS. I.e., out-adjacency list is used for the exploration of DFS in the forward direction and the in-adjacency list is used for finding parent of a node during the backtracking step. We provide the details below. Similar to our algorithm for undirected graphs, in the forward direction, we scan the out-neighbor list of w to find the next white neighbor and proceed. Once the out-neighbor list of w is fully processed, we need to backtrack from w . Towards that we first have to identify w 's parent. In order to do so we use the following lemma whose proof follows along the same lines as the Lemma 1 above. Hence we omit the proof.

Lemma 2. *Suppose w is a node that just became black. Then its parent p is the unique vertex in w 's in-adjacency list which is (a) gray and (b) whose current out-adjacency list has w in the first position.*

Once we figure out w 's parent p , DFS backtracks to p , finds the next white neighbor (as done in the forward step) and continues until all the vertices are exhausted. It is clear that this procedure performs lex-DFS on a directed graph G correctly in linear time, and this completes the proof of Theorem 1(a).

2.3. General-DFS using $n + O(\lg n)$ bits–proof of Theorem 1(b) for undirected graphs

To improve the space further, we replace the color array C with a bit array $visited[1, \dots, n]$ which stores a 0 for an unvisited vertex (white), and a 1 for a visited vertex (gray or black).

First we need a test similar to that in the statement of Lemma 1 without the distinction of gray and black vertices to find the parent of a node. Due to the invariant we have maintained, every internal vertex of the DFS tree will point to (i.e. have as first element in its list) its last child. So the nodes that could potentially have a node w in its first position are its parent, and any leaf vertex. Hence we modify the forward step in the following way.

Whenever we visit an unvisited vertex v for the first time from another vertex u (hence, u is the parent of v in the DFS tree and u 's list has v in the first position), we, as before, mark v as visited and in addition to that, we rotate v 's list to bring u to the front (during this rotation, we do not mark any intermediate nodes as visited). Then we continue as before (by finding the first unvisited vertex and bringing it to the front) in the forward step. Now the following invariants are easy to see and are useful.

Invariants: During the exploration of DFS, in the (partial) DFS tree

1. any internal vertex has the first element in its list as its current last child; and
2. for any leaf vertex of the DFS tree, the first element in its list is its parent.

The first invariant is easy to see as we always keep the current explored vertex (child) as the first element in the list. For leaves, the first time we encounter them, we make its parent as the first element in the forward direction. Then we discover that it has no unvisited vertices in its list, and so we make a full rotation and bring the parent to the front again. The following lemma provides a test to find the parent of a node.

Lemma 3. *Let w be a node that has just become black. Then its parent p is the **first** vertex x in w 's list whose current adjacency list has w in the first position.*

Proof. From the invariants we observed, the nodes that can potentially have w in the first position of their lists are its parent and its children that happen to be leaves. But in w 's list, as we began the exploration of its neighbors starting from its parent, its parent will appear first before its children. Hence the first node in w 's list which has w in the first position must be its parent. \square

Once we backtrack to p , we find the next white vertex, and continue until all the vertices of G are explored. Overall this procedure takes linear time. As we rotate the list to bring the parent of a node, before exploring its white neighbors, we are not guaranteed to explore the first white vertex in its original list, and hence we loose the lexicographic property. We provide our DFS algorithm for directed graphs below.

2.4. General-DFS using $n + O(\lg n)$ bits–proof of Theorem 1(b) for directed graphs

For performing DFS in directed graphs using $n + O(\lg n)$ bits, we don't even need to apply the modifications as we did for the undirected graphs during the forward direction, and we can essentially use the same forward step idea as used for lex-DFS in undirected graphs of Section 2.1. We provide the details below. When we arrive at a previously unvisited vertex v from the vertex u (hence u is the parent of v in the DFS tree), we rotate the in-neighbor list of v to bring u to the front and u stays there during the entire course of the exploration. Thus we maintain the invariant that for any visited node v ,

the first element in its in-neighbor list is its parent in the DFS tree. Now the algorithm scans v 's adjacency list to find its unvisited neighbor. (i) If we find such a vertex, say w , then we rotate v 's list to make w as the first element, and mark w visited. (ii) If we do not find any such unvisited neighbor of v , then DFS needs to *backtrack* to its parent. From the invariant we maintain in the in-neighbor list of every visited vertex, this is easy. All we need to do is to see the first entry in v 's in-neighbor list to retrieve its parent u and then continue from u . Overall this procedure takes linear time. Upon closer inspection, it can be seen that, as we are not modifying the ordering of the vertices in the out-neighbor lists in the forward direction (in contrast with the undirected graph algorithm of Section 2.3), this procedure actually traverses the directed graph G in lex-DFS ordering. This completes the proof of Theorem 1(b).

2.5. General-DFS using $O(\lg n)$ bits—proof of Theorem 1(c) for undirected graphs

Now to decrease the space to $O(\lg n)$, we dispense with the color/visited array, and give tests to determine white, gray and black vertices. For now, assume that we can determine the color of a vertex. The forward step is almost the same as before except performing the update in the color array. I.e., whenever we visit a white vertex v for the first time from another vertex u (hence u is the parent of v), we rotate v 's list to bring u to the front. Then we continue to find the first white vertex to explore. We maintain the following invariants. (i) any gray vertex has the first element in its list as its last child in the (partial) DFS tree; (ii) any black vertex has its parent as the first element in its list. We also store the depth of the current node (i.e., the last gray vertex in the gray path of the DFS tree) in a variable d , which is incremented by 1 every time we discover a white vertex and decremented by 1 whenever we backtrack. We maintain the maximum depth the DFS has attained using a variable max . At a generic step during the execution of the algorithm, assume that we are at a vertex x 's list, let p be x 's parent and let y be a vertex in x 's list. We need to determine the color of y and continue the DFS based on the color of y . We use the following characterization.

Lemma 4. *Suppose the DFS has explored starting from a source vertex s , up to a vertex x at level d . Let p be x 's parent. Note that both s and x are gray in the normal coloring procedure. Let max be the maximum level of any vertex in the partial DFS exploration. Let y be a vertex in x 's list. Then,*

1. y is gray (i.e., (x, y) is a back edge, and y is an ancestor of x) if and only if we can reach y from s by following through the gray child (which is at the front of a gray node's list) path in at most d steps.
2. y is black (i.e., (x, y) is a back edge, and x is an ancestor of y) if and only if
 - there is a path P of length at most $(max - d)$ from y to x (obtained by following through the first elements of the lists of every vertex in the path, starting from y), and
 - let z be the node before x in the path P . The node z appears after p in x 's list.
3. y is white if y is not gray or black.

Proof. The test for gray and white vertices is easy to see. The vertex y is black implies that y is a descendant of x in the partially explored DFS tree. This means that there is a path of length at most $(max - d)$ (obtained by following the parent which is in the first element of the adjacency list) from x to y through an already explored child z . By the way we process x 's list, we first bring the parent to the front of the list, and then explore the nodes in sequence, and hence z , the explored neighbor of x must appear after p in x 's list. Conversely, the unexplored neighbors of x appear before p in x 's list. \square

Now, if we use the above claim to test for colors of vertices, testing for gray takes at most d steps. Testing for black takes at most $(max - d)$ steps to find the path, and at most $deg(x)$ steps to determine whether p appears before. Thus for each vertex in x 's list, we spend time proportional to $max + deg(x)$. So, the overall runtime of the algorithm is $\sum_{v \in V} deg(v)(deg(v) + \ell) = O(m^2/n + m\ell)$, where ℓ is the maximum depth of DFS tree. Maintaining the invariants for the gray and black vertices are also straightforward. We provide the details of our log-space algorithm for directed graphs below.

2.6. General-DFS using $O(\lg n)$ bits—proof of Theorem 1(c) for directed graphs

We describe our $O(\lg n)$ bits algorithm for directed graphs in the rotate model. More specifically, we give a DFS algorithm to output all vertices reachable by a directed path from the source vertex s . If we assume that s can reach all vertices, we get to output all vertices. In the preprocessing step, the algorithm spends $O(m)$ time to bring the minimum valued neighbor (denote it by min) in the out-neighbor list of every vertex by rotation (hence we lose the lexicographic DFS property). For now assume that we can determine the color of a vertex. Given this, in the forward direction, when DFS arrives at a previously unvisited vertex v from the vertex u (hence u is the parent of v in the DFS tree), we rotate the in-neighbor list of v to bring u to the front and u stays there during the entire course of the exploration. Also in u 's out-neighbor list, v is made the first location. Hence we maintain the following invariants.

Invariants: During the exploration of DFS, in the (partial) DFS tree




Fig. 2. Illustration of the different cases of the possible positions of the vertex y when DFS considers the directed edge (x, y) at some intermediate step. Suppose the root of the DFS tree is the vertex s and the curvy path starting from s and going straight below through x is the current gray path in the DFS tree. Intuitively all the vertices on the left hand side of the path are black, and right hand side are white and yet to be explored. From left to the right are cases when (x, y) is (a) back edge, (b) cross edge, (c) forward edge, and (d) tree edge.

1. gray vertices have their current last child in the first location of their out-neighbor lists;
2. all the visited (i.e., gray and black) vertices have their parent in the first location of their in-neighbor lists.

We also keep track of the depth of the current node in a variable d , which, as before, is incremented by 1 every time DFS visits a white vertex and decremented by 1 whenever DFS backtracks. We also store the maximum depth the DFS tree has attained so far in a variable max . At a generic step during the execution of the algorithm, assume that we are at a vertex x 's list, let p be x 's parent (which can be found from the way x is visited by a forward or a backtracking step using the invariants being maintained) and let y be a vertex in x 's list. We need to determine the color of y and continue the DFS based on the color of y and maintain the invariants. We use the following characterization.

Lemma 5. *Suppose the DFS has explored starting from a source vertex s up to a vertex x at level d . Let p be x 's parent. Note that both s and x are gray in the normal coloring procedure. Let max be the maximum level of any vertex in the partial DFS exploration. Let y be a vertex in x 's list. Then,*

1. y is gray (i.e., (x, y) is a back edge and y is an ancestor of x) if and only if we can reach y from s following the gray child (which is in the first location of each of the out-neighbor lists of gray nodes) path in at most d steps.
2. y is black if and only if any of the following happens.
 - There is a path P of length at most $(max - d)$ from y to x obtained by following the first elements of the in-neighbor lists of every vertex in the path P starting from y . This happens when (x, y) is a forward edge, and x is an ancestor of y .
 - There is a path P of length at most max from y to a gray vertex $z \neq x$ (obtained by following through the first elements of the in-neighbor lists of every vertex starting from y) which is the first gray vertex in the path. Let c be the node before z in the path P , then c must appear after min in z 's list (this happens when (x, y) is a cross edge).
3. The vertex y is white if it is not black or gray (i.e., (x, y) is the next tree edge with x being the parent of y in the DFS tree).

Proof. See Fig. 2 for a picture of all the cases. The test for gray and white vertices is easy to see.

From a vertex x , there could be two types of outgoing edges to a black vertex y . When (x, y) is a forward edge, y is a descendant of x and hence there must exist a path P of length at most $(max - d)$ (obtained by following the parent which is in the first location of the in-neighbor list of every vertex in P , starting from y) from y to x through an already explored child t of x . In the other case, when (x, y) is a cross edge, y has already been discovered and explored completely before DFS reaches to x . Hence there must exist a gray ancestor z of x (z could be x) such that y belongs to the subtree rooted at z in the DFS tree. Thus, from y 's in-neighbor list if we follow the path starting with y 's parent for at most max steps, we must hit the gray path and the first vertex we come across is z . Let c be the node before z in the path. By the way we process z 's out-neighbor list, we first bring the min to the front of the list, and then explore the other neighbor nodes in sequence, and hence c , the explored neighbor of z must appear after min in z 's list.

For the converse, suppose y is a white vertex. Either we never reach a gray vertex in max steps (and we will correctly determine its color in this case) or we reach x or x 's ancestor z from y following through the (spurious) first vertices of the in-neighbor list of a white vertex y . Note that the parent of a white vertex is white or gray and it can never be black. Hence z 's child in the path is white. Hence that child will appear before min in z 's list. \square

Given the above lemma, if y turns out to be white, the edge (x, y) is added to the DFS tree, and y now becomes the current gray vertex. Note that maintaining the invariants are straightforward. Also, when any vertex v has been completely explored by DFS, we retrieve its parent from the in-neighbor list to complete the backtracking step. This procedure is

continued until DFS comes back to the source vertex s . We stop at this point. This is because, note that, our proof breaks down in the case when DFS in a directed graph produces a forest and some cross edges go across these trees. In that case, if we follow the path starting from y 's parent, we would reach the root of the DFS tree containing y and this is different from the tree where x belongs to. As we cannot maintain information regarding all such roots of these previously explored DFS trees, we might spuriously conclude that y is unvisited even though it is not the case. Thus our algorithm produces the DFS tree containing only the vertices reachable from the source vertex s via some directed path in G . We leave open the case for designing such logspace algorithm for the general directed graphs.

To analyze the running time of our algorithm observe that testing for gray takes at most d steps. Testing for black takes, in the worst case, at most max steps to find the path, and at each step of the path, we take d time to test whether the new node is gray. Once we reach a gray vertex, we spend at most $deg(z)$ steps to determine whether c appears before min . Thus for each vertex in x 's list, we spend time proportional to $d + (d \cdot max) + deg(z)$ time. As z (which is independent of x) can have degree at most n . Thus, the overall runtime of the algorithm is $\sum_{v \in V} deg(v)(d + d\ell + n)$ which is $O(m(n + (1 + \ell)\ell))$ which is $O(m(n + \ell^2))$, where ℓ is the maximum depth of DFS tree.

3. BFS algorithms in the rotate model

3.1. BFS using $n + O(\lg n)$ bits—proof of Theorem 2(a)

It is well-known that BFS actually computes the shortest path lengths in unweighted undirected or directed graph G from a given source vertex $s \in V$ to every vertex $v \in V$ that is reachable from s . I.e., if a vertex v belongs to the d -th level in the BFS tree (assuming the root s is at zero-th level), then we know that the length of the shortest path from s to v is d . We use this crucially to design our BFS algorithms. We use a bit array $visited[1, \dots, n]$ that stores a 0 for an unvisited vertex, and 1 for a visited vertex. We also maintain a counter $dist$ which stores the level of the vertex that is currently being explored in the BFS algorithm.

We start by setting $visited[s] = 1$, and initializing the counter $dist$ to 0. At the next step, for every unvisited neighbor v of s , we rotate their adjacency list so that s appears as the first element in v 's list, set $visited[v] = 1$, and output v . This step ensures that for each visited vertex, its parent is at the front of its adjacency list. We refer to this front element in the adjacency list of a visited vertex as its *parent pointer*. (Also, once we set the parent pointer for a vertex, we will not rotate its adjacency list in the remaining part of the algorithm.) Once the root s 's list is fully processed as above, the $dist$ is incremented to 1. The next step in the algorithm is to find all the vertices in the first level and mark all their unvisited neighbors as visited. As we haven't stored these vertices (in the first level), the challenge is to find them first. We use the following claim, to find the level number of a visited vertex. The proof easily follows from the fact that the parent pointers are set for all the visited vertices, and that all the ancestors of a visited vertex are also visited.

Claim 1. *If the BFS has been explored till distance d from the source, then for any $k \leq d$, a vertex x marked visited is in level k if and only if we can reach the source s in exactly k steps by following through their parent pointers. Thus determining if a visited vertex is in level d takes at most d steps.*

So now we continue the BFS by scanning through the vertices, finding those vertices in level d (using the above claim by spending d steps for each vertex), and marking their unvisited neighbors visited, and making in their adjacency lists, their parent vertex as the first element, and incrementing $dist$. We stop our algorithm when we discover no new unvisited vertex while exploring any level. The correctness of the procedure and the space used by the algorithm are clear. To analyze the runtime, note that the time spent at level d is $nd + \sum_{i \in V(d)} deg(i)$ where $V(d)$ is the set of vertices in level d and $deg(i)$ is the degree of vertex i . Summing over all levels, we get a runtime $O(m + n\ell^2)$, where ℓ is the depth of the BFS tree. To handle directed graphs, we follow the outneighbor list as we go down, and we set up the parent at the first position in the in-neighbor list of every vertex v . To verify if v belongs to the d -th level, we take d steps from v by following the parent pointers in the in-neighbor lists of the (visited) vertices along the path, and check if we reach s at the end. This completes the proof of Theorem 2(a).

3.2. BFS using $O(\lg n)$ bits—proof of Theorem 2(b)

To reduce the space to $O(\lg n)$ bits, we dispense with the color array and explain how to determine visited and unvisited vertices. Assume that we can determine this in constant time. Our first observation is that Claim 1 is true even for unvisited vertices even though the first vertex in the adjacency list of unvisited vertices can be an arbitrary vertex (not necessarily referring to their parent in the BFS tree). However, we know (by the property of BFS) that no unvisited vertex is adjacent to a vertex in level less than d , and hence they can not reach s by a path at most d . Using the same argument, we can show that

Claim 2. *If vertices up to level d have been explored and visited, then a vertex x is a visited vertex if and only if by following through the parent pointers, x can reach s by a path of length at most d . Furthermore, a vertex is in level d if and only if we can reach s by a path of length exactly d by following through the parent pointers.*

Thus, to check whether a vertex is visited, we spend $O(d)$ time when exploring vertices at level d instead of $O(1)$ time when the *visited* array was stored explicitly. Hence, the total time spent at level d is $O(nd + d \sum_{i \in V(d)} \text{deg}(i))$, where the first term gives the time to find all the visited vertices at level d , and the second term gives the time to explore those vertices (i.e., going through their neighbor lists, identifying the unvisited neighbors and setting their parent pointers). Summing over all levels, we get a total runtime of $O(n\ell^2 + m\ell)$. Note that this algorithm works only when the input undirected graph is connected as Claim 2 breaks down if there are more than one component. The modifications to handle the directed graphs are similar to those for the directed graph BFS algorithm. This proves Theorem 2(b).

4. Simulation of algorithms for rotate model in the implicit model

The following result captures the overhead incurred while simulating any rotate model algorithm in the implicit model. Most of our algorithms in the implicit model use these simulations often with some enhancements and tricks to obtain better running time bounds for some specific problems.

Theorem 6. *Let D be the maximum degree of a graph G . Then any algorithm running in $t(m, n)$ time in the rotate model can be simulated in the implicit model in (i) $O(D \cdot t(m, n))$ time when G is given in an adjacency list, and (ii) $O(\lg D \cdot t(m, n))$ time when G is given in an adjacency array. Furthermore, let $r_v(m, n)$ denote the number of rotations made in v 's (whose degree is d_v) list, and $f(m, n)$ be the remaining number of operations. Then any algorithm running in $t(m, n) = \sum_{v \in V} r_v(m, n) + f(m, n)$ time in the rotate model can be simulated in the implicit model in (i) $O(\sum_{v \in V} r_v(m, n) \cdot d_v + f(m, n))$ time when G is given in an adjacency list, and (ii) $O(\sum_{v \in V} r_v(m, n) \lg d_v + f(m, n))$ time when G is given in an adjacency array.*

Proof. We can implement a single rotate operation that moves the head of the list by one position (which is assumed to be a unit-cost operation in the rotate model) in the implicit model by moving all the elements in the adjacency list circularly. If d_v is the degree of a vertex v , then performing a rotation of the adjacency list of v can be implemented in $O(d_v)$ time in the implicit model. Thus, if we have an algorithm in the rotate model that takes $t(m, n)$ time, then it can be implemented in $O(D \cdot t(m, n))$ time in the implicit model, where D is the maximum degree of the graph. One can get a better runtime by analyzing the algorithm for the rotate model more carefully. In particular, if the runtime of the algorithm in the rotate model can be expressed as $r(m, n) + f(m, n)$, where $r(m, n)$ is the number of rotations performed and $f(m, n)$ is the remaining number of operations, then the algorithm can be implemented in the implicit model in $O(D \cdot r(m, n)) + f(m, n)$ time. Furthermore, if $r(m, n) \leq \sum_{v \in V} r_v(m, n)$ where $r_v(m, n)$ is the number of rotations made in v 's list, then the runtime of the algorithm in the implicit model can be bounded by $O(\sum_{v \in V} r_v(m, n) \cdot d_v + f(m, n))$.

If the input graph is given in the adjacency array representation and if the ordering of the elements in the adjacency array is allowed to be changed, then one can simulate the rotate operation even faster. The main idea is to simulate the algorithm for the rotate model after sorting the adjacency arrays of each vertex. Using an in-place linear time radix sorting algorithm [60], sorting all the adjacency arrays can be done in $O(m)$ time. Next, we note that in the rotate model, the head points to an element in an arbitrary position (called the front element) in the adjacency list of any vertex and the unit operation moves it one position. Thus, it is enough to show how to access the next element in sorted order. We maintain the following invariant: if the first element is the i -th element in sorted order, then the adjacency array consists of the sorted array of all the adjacent vertices, with the first element swapped with the i -th element. To bring the $(i + 1)$ -st element to the front, we first perform a binary search for the i -th element which is in the first position in the 'almost sorted' adjacency array to find the position i , and move the elements appropriately to maintain the invariant (with the $(i + 1)$ -st element at the front). This takes $O(\lg d_v)$ time to simulate the rotation of the adjacency array of a vertex v with degree d_v . Thus, if we have an algorithm in the rotate model that takes $t(m, n)$ time, then it can be implemented in $O(\lg D \cdot t(m, n))$ time in the implicit model. Moreover, if the runtime of the algorithm in the rotate model can be expressed as $\sum_{v \in V} r_v(m, n) + f(m, n)$, where $r_v(m, n)$ is an upper bound on the number of rotations performed on vertex v and $f(m, n)$ is the remaining number of operations, then the algorithm can be implemented in the implicit model in $O(\sum_{v \in V} r_v(m, n) \lg d_v + f(m, n))$ time. \square

5. DFS algorithms in the implicit model

To obtain a lex-DFS algorithm, we implement the $O(\lg n)$ -bit DFS algorithm in the rotate model, described in Section 2.5, with a simple modification. First, note that in this algorithm (in the rotate model), we bring the parent of a vertex to the front of its adjacency list (by performing rotations) when we visit a vertex for the first time. Subsequently, we explore the remaining neighbors of the vertex in the left-to-right order. Thus, for each vertex, if its parent in the DFS were at the beginning of its adjacency list, then this algorithm would result in a lex-DFS algorithm. Now, to implement this algorithm in the implicit model, whenever we need to bring the parent to the front, we simply bring it to the front without changing the order of the other neighbors. Subsequently, we simulate each rotation by moving all the elements in the adjacency list circularly. As mentioned in Section 4, this results in an algorithm whose running time is $O(\sum_{v \in V} d_v(d_v + \ell) \cdot d_v) = O(m^3/n^2 + \ell m^2/n)$ if the graph is given in an adjacency list and in $O(\sum_{v \in V} d_v(d_v + \ell) \cdot \lg d_v) = O(m^2(\lg n)/n + m\ell \lg n)$ when the graph is given in the form an adjacency array. This proves Theorem 3(a) for undirected graphs. The results for the directed case follow from simulating the corresponding results for the directed graphs.

To prove the result mentioned in Theorem 3(b), we implement the linear-time DFS algorithm of Theorem 1 for the rotate model that uses $n + O(\lg n)$ bits. This results in an algorithm that runs in $O(\sum_{v \in V} d_v^2 + n) = O(m^2/n)$ time (or in $O(\sum_{v \in V} d_v \lg d_v + n) = O(m \lg m + n)$ time, when the graph is given as an adjacency array representation), using $n + O(\lg n)$ bits. We reduce the space usage of the algorithm to $O(\lg n)$ bits by encoding the visited/unvisited bit for each vertex with degree at least 2 within its adjacency list (and not maintaining this bit for degree-1 vertices). We describe the details below.

Whenever a node is visited for the first time in the algorithm for the rotated list model, we bring its parent to the front of its adjacency list. In the remaining part of the algorithm, we process each of its other adjacent vertices while rotating the adjacency list, until the parent comes to the front again. Thus, for each vertex v with degree d_v , we need to rotate v 's adjacency list $O(d_v)$ times. In the implicit model, we also bring the parent to the front when a vertex is visited for the first time, for any vertex with degree at least 3. We use the second and third elements in the adjacency list to encode the visited/unvisited bit. But instead of rotating the adjacency list circularly, we simply scan through the adjacency list from left to right everytime we need to find the next unvisited vertex in its adjacency list. This requires $O(d_v)$ time for a vertex v with degree d_v . We show how to handle vertices with degree at most 2 separately.

As before, we can deal with the degree-1 vertices without encoding visited/unvisited bit as we encounter those vertices only once during the algorithm. For degree-2 vertices, we initially (at preprocessing stage) encode the bit 0 using the two elements in their adjacency arrays - to indicate that they are unvisited. When a degree-2 vertex is visited for the first time from a neighbor x , we move to its other neighbor - continuing the process as long as we encounter degree-2 vertices until we reach a vertex y with degree at least 3. If y is already visited, then we output the path consisting of all the degree-2 vertices and backtrack to x . If y is not visited yet, then we output the path upto y , and continue the search from y , and after marking y as visited. In both the cases, we also mark all the degree-2 nodes as visited (by swapping the two elements in each of their adjacency arrays).

During the preprocessing, for each vertex with degree at least 3, we ensure that the second and third elements in its adjacency list encode the bit 0 (to mark it unvisited). We maintain the invariant that for any vertex with degree at least 3, as long as it is not visited, the second and third elements in its adjacency array encode the bit 0; and after the vertex is visited, its parent (in the DFS tree) is at the front of its adjacency array, and the second and third elements in its adjacency array encode the bit 1. Thus, when we visit a node v with degree at least 3 for the first time, we bring its parent to the front, and then swap the second and third elements in the adjacency list, if needed, to mark it as visited. The total running time of this algorithm is bounded by $\sum_{v \in V} d_v^2 = O(m^2/n)$.

We can implement the above DFS algorithm even faster when the input graph is given in an adjacency array representation. We deal with vertices with degree at most 2 exactly as before. For a vertex v with degree at least 3, we bring its parent to the front and swap the second and third elements to mark the node as visited (as before) whenever v is visited for the first time. We then sort the remaining elements, if any, in the adjacency array, in-place (using the linear-time in-place radix sort algorithm [60]), and implement the rotations on the remaining part of the array as described in Section 4. The total running time of this algorithm is bounded by $\sum_{v \in V} d_v \lg d_v = O(m \lg m + n)$. This completes the proof of Theorem 3(b).

6. BFS algorithms in the implicit model

Before getting into the technical details, we first outline the main ideas involved in proving Theorem 4. One can simulate the BFS algorithm of Theorem 2(b) (for the rotate model) in the implicit model using the simulation described in Section 4. Since these BFS algorithms scan through each of the adjacency lists/arrays at most twice during the algorithm, there won't be any slowdown in the runtime. This results in an algorithm running in $O(m\ell + n\ell^2)$ time, using $O(\lg n)$ bits. To improve the running time, we simulate the algorithm of Theorem 2(a) using the trick of encoding the *visited* bit of each vertex in its adjacency list instead of storing the *visited* array explicitly. This requires special attention to degree-1 and degree-2 vertices along with few other technical issues which are dealt in the proof given next.

Proof. Here we give the full proof of Theorem 4. In particular, we provide all the details of the case when the degree of each vertex is at least 3, and in that case, we show that we can implement the BFS algorithm using 4 colors of [14], by encoding the 4 color of a vertex using the first three elements in its adjacency list, resulting in an algorithm that takes $O(m + n\ell)$ time. Moreover, when the degree of every vertex is at least $2 \lg n + 3$, then we show that the above algorithm can be implemented more efficiently, resulting in an algorithm that takes $O(m)$ time. Details follow.

To improve the running time further, we simulate the algorithm of Theorem 2(a). But instead of storing the *visited* array explicitly, we encode the *visited* bit of each vertex in its adjacency list, as explained below, resulting in an algorithm that takes $O(m + n\ell^2)$ time, using $O(\lg n)$ bits. For a vertex x with degree at least 3, we encode its *visited* bit using the second and third elements in its adjacency list. To set the parent pointer for x (when it is visited for the first time), we bring its parent to the front, and move the second and third elements, if necessary, to encode the *visited* bit. We now describe how to deal with the degree-1 and degree-2 vertices.

First, observe that in the original BFS algorithm, we can simply output any degree-1 vertex, when it is visited for the first time. Thus, we need not store the *visited* bit for degree-1 vertices. For degree-2 vertices, we encode the *visited* bit using the two neighbors. We do not bring its parent to the front, as we do for other vertices. Whenever we need to check whether a visited degree-2 vertex is at depth d , we follow parent pointers from both the neighbors - if one of them does not exist, then the other one is its parent - for a distance of length at most d . (While following the parent pointers from

a vertex to the root, it is easy to find its parent - since there is only one alternative to follow.) Thus, the first part of the theorem follows from Theorem 2(a), with the above modification.

To improve the runtime further, we implement the BFS algorithm using 4 colors by [14], where all the unvisited vertices are white, all the visited, yet unvisited vertices of the two consecutive layers of BFS are gray1 and gray2 respectively, and all the vertices which are completely explored are black. Suppose the degree of every vertex in the given graph is at least three. In this case, we can encode the color of any vertex by permuting the first three elements in its adjacency array appropriately. This enables us to retrieve (or modify) the color of any vertex in constant time by reading (or modifying) the permuted order of the first elements in its adjacency array.

Since we haven't stored the gray1 or gray2 vertices in a queue (as in the standard BFS), we scan through the entire vertex set (in the increasing order of their labels), and when we find any vertex v colored gray1, we color all its white neighbors with gray2, and color v itself with black. We call this an *exploration phase*. The time taken for each exploration phase is equal to the sum of the degrees of all the gray1 vertices at the beginning of the phase, plus $O(n)$. At the end of each exploration phase, we change the colors of all gray2 vertices to gray1 and also output them, using an additional $O(n)$ time. We call this a *consolidation phase*. We need to repeat the exploration and consolidation phases for ℓ times before all the nodes are colored black, where ℓ is the height of the BFS tree. Thus the overall time taken by this procedure can be bounded by $O(m + n\ell)$. This proves the second part of the theorem.

If every vertex has degree at least $2\lg n + 3$, then we can perform BFS in $O(m)$ time (note that $m \geq n\lg n$ in this case) - by encoding the current set of gray1 vertices as a linked list by using $2\lg n$ vertices to encode (a pointer to) the next vertex in the list. The time to read all the gray1 vertices in a phase when there are k vertices colored gray1 becomes $O(k\lg n)$ instead of $O(n)$. This results in $O(m + n\lg n)$ time which is $O(m)$. This proves the third part of the theorem. \square

7. Minimum spanning tree

In this section, we start by giving an in-place implementation of the Prim's algorithm [38] to find a minimum spanning tree of a given weighted undirected graph in the rotate model. Here we are given a weight function $w : E \rightarrow Z$. We also assume that the weights of non-edges are ∞ and that the weights can be represented using $O(\lg n)$ bits. The input representation also changes slightly to accommodate these weights. Now each element of the circular linked list has three fields, (a) the vertex label, (b) the weight, and (c) the pointer to the next vertex respectively. In what follows, when we talk about computing a minimum spanning tree, what we mean is reporting the edges of such a tree. Our result is the following,

Theorem 7. *A minimum spanning tree of a given undirected weighted graph G can be found using $O(\lg n)$ bits and in $O(mn)$ time in the rotate model.*

Proof. Our rotate model algorithm basically mimics Prim's algorithm with a few tweaks. Prim's algorithm starts with initializing a set S with a vertex s . For every vertex v not in S , it finds and maintains $d[v] = \min\{w(v, x) : x \in S\}$ and $\pi[v] = x$ where $w(v, x)$ is the minimum among $\{w(v, y) : y \in S\}$. Then it repeatedly deletes the vertex with the smallest d value from $V - S$ adding it to S . Then the d values are updated by looking at the neighbors of the newly added vertex. To implement this procedure using just extra $O(\lg n)$ bits of space, first we need to find a way to mark/unmark a vertex v if it has been taken into S without using n bits explicitly. The way we do this is as follows. In the preprocessing step, the algorithm spends $O(m)$ time to bring the minimum valued neighbor (denote it by *min*) in the neighbor list of every vertex v by rotation. Subsequently we would attach the following meaning with the position of *min* in the list of any vertex v . If the first element in the list of any vertex v is *min*, this means that v is not taken into S so far during the execution of the algorithm; otherwise it belongs to S . This way we can store the information regarding the status of any vertex v without using any extra space but actually figuring out this information takes time proportional to the degree of v . Note that for vertices having degree one, we cannot determine exactly its status correctly by this method. But a simple fact we can use here. If a vertex z has degree one (say its neighbor is y), then the edge (y, z) is always a part of the minimum spanning tree. Hence, after the preprocessing step, we can output all such edges at once and embark on executing the rest of the algorithm.

The algorithm initializes the set S with the starting vertex s , and goes to its list to give a rotation so that *min* does not stay in the first location in s 's list. We call this step as the *marking* step. This is followed by finding the smallest weight neighbor (say u) of s . According to Prim's algorithm, u should now move to S . We achieve the same by marking u i.e., going to u 's list to give a rotation to move *min* from the first location to indicate that u belongs to S now and subsequently continue to in u 's list find its smallest weight neighbor and repeat. Thus at any generic step of the algorithm, we go over the list of unmarked vertices (i.e., those vertices having *min* at the first position of their respective lists) and collect the minimum weight vertex (say t), and t is then marked and we continue until all the vertices are marked.

Clearly this method returns all the correct edges of a minimum spanning tree of G . Space bound of this algorithm is easily seen to be $O(\lg n)$ bits for keeping a few variables. Total time spent by algorithm can be bounded by $O(mn)$ where at the preprocessing step, it spends $O(m)$ time and after the preprocessing, for reporting each edges of the minimum spanning tree, in the worst case, the algorithm spends $O(m)$ time; hence $O(mn)$ is the bound on the running time. \square

As mentioned in Section 4, simulating this algorithm in the implicit model would result in an algorithm having running time $O(mn \cdot d_v) = O(mn^2)$ if the graph is given in an adjacency list and in $O(mn \cdot \lg d_v) = O(mn \lg n)$ when the graph is represented in an adjacency array. Hence, we have the following,

Theorem 8. *In the implicit model a minimum spanning tree of a given undirected weighted graph G can be found using $O(\lg n)$ bits and*

1. $O(mn^2)$ time if the graph is given in an adjacency list, and
2. $O(mn \lg n)$ time when the graph is represented in an adjacency array.

8. Consequences

In this section, we provide some applications/consequences of our results that we derived in the earlier sections.

8.1. Improved algorithm for DFS and applications in ROM

In this section, we show, using a little more space, how to simulate any rotate model algorithm in the read-only model. This results in improved space-efficient algorithms for various fundamental graph problems in the read-only model. Observe that, the only modification of the input that we do in our rotate model algorithm is to make the head pointer point to an arbitrary element in the adjacency list (instead of a fixed element) at various times. To simulate this in read-only memory, we can simply maintain a pointer in each of the lists in the adjacency list. The resources required to store and update such pointers is proven in the following lemma [30]. We provide the proof here for completeness.

Lemma 6 ([30]). *Given the adjacency list representation of a directed or an undirected graph G on n vertices with m edges, using $O(m)$ time, one can construct an auxiliary structure of size $O(n \lg(m/n))$ bits that can store a “pointer” into an arbitrary index within the adjacency list of each vertex. Also, updating any of these pointers (within the adjacency list) takes $O(1)$ time.*

Proof. We first scan the adjacency list of each vertex and construct a bitvector B as follows: starting with an empty bitvector B , for $1 \leq i \leq n$, if d_i is the length of the adjacency array of vertex v_i (i.e., its degree), then we append the string $0^{\lceil \lg d_i \rceil - 1} 1$ to B . The length of B is $\sum_{i=1}^n \lceil \lg d_i \rceil$, which is bounded by $O(n \lg(m/n))$. We construct auxiliary structures to support *select* queries on B in constant time [71]. We now construct another bitvector P of the same size as B , which stores the pointers into the adjacency array of each vertex. The pointer into the adjacency array of vertex v_i is stored using the $\lceil \lg d_i \rceil$ bits in P from position $\text{select}(i-1, B) + 1$ to position $\text{select}(i, B)$, where $\text{select}(0, B)$ is defined to be 0. Now, using *select* operations on B and using constant time word-level read/write operations, one can access and/or modify these pointers in constant time. \square

To actually get to the element in the list, we need the graph to be represented as what is referred as adjacency array [52]. Here given an index in the list of a vertex, we can access the (adjacent) vertex in that position of the vertex's adjacency list in constant time. Now if we simulate our rotate model algorithm of Theorem 1 in read-only memory using the auxiliary structure as stated in Lemma 6 as additional storage, then we obtain,

Theorem 9. *A DFS traversal of an undirected or a directed graph G , represented by an adjacency array, on n vertices and m edges can be performed in $O(m+n)$ time using $O(n \lg(m/n))$ bits, in the read-only model.*

The above result improves the DFS tradeoff result of Elmasry et al. [52] for relatively sparse graphs in the read-only model. In particular, they showed the following,

Theorem 10 ([52]). *For every function $t: \mathbb{N} \rightarrow \mathbb{N}$ such that $t(n)$ can be computed within the resource bound of this theorem (e.g., in $O(n)$ time using $O(n)$ bits), the vertices of a directed or undirected graph G , represented by adjacency arrays, with n vertices and m edges can be visited in depth first order in $O((m+n)t(n))$ time with $O(n + n \frac{\lg \lg n}{t(n)})$ bits.*

Thus to achieve $O(m+n)$ time for DFS, their algorithm (Theorem 10) uses $O(n \lg \lg n)$ bits. This is $\Omega(n \lg(m/n))$ for all values of m where $m = O(n \lg n)$. Banerjee et al. [14] and Kammer et al. [63] recently provided another DFS implementation taking $O(m+n)$ bits of space and runs in $O(m+n)$ time. Note that, Theorem 9 improves the space bound of the above mentioned DFS implementations from $O(m+n)$ space to $O(n \lg(m/n))$, while maintaining the same linear running time. Chakraborty et al. [30] obtained similar results by a slightly different technique. In what follows, we show that using Theorem 9, we can improve the space bounds of some of the classical applications of DFS in read-only model.

To illustrate this, note that, one of the many classical applications of DFS (see [38]) include (i) topological sorting of the vertices of a directed acyclic graph [65], (ii) producing a sparse (having $O(n)$ edges) spanning biconnected subgraph of a

undirected biconnected graph G [51], and (iii) given an undirected 2-edge-connected graph $G = (V, E)$, to orient each edge $(u, v) \in E$ as an arc (u, v) or (v, u) to obtain $D = (V, A)$ such that D becomes strongly connected. For all of these problems, classical algorithms [65,51] take linear time and $O(n \lg n)$ bits of space. Recently, Banerjee et al. [14] showed the following,

Theorem 11. [14] *In the read-only model, if the DFS of G on n vertices and m edges, can be performed in $t(m, n)$ time using $s(m, n)$, where $s(m, n) = \Omega(n)$, bits of space, then using $O(s(m, n))$ bits and in $O(t(m, n))$ time, we can output*

1. the vertices of a directed acyclic graph in topologically sorted order,
2. the edges of a sparse spanning biconnected subgraph of a undirected biconnected graph G , and
3. a strongly connected orientation of a undirected 2-edge-connected graph G .

Now plugging the improved DFS algorithm of Theorem 9 in the above theorem, we obtain for all the applications an improved (over the classical implementation) space-efficient implementations taking $O(m + n)$ time and $O(n \lg(m/n))$ bits of space.

8.2. Space-efficient approximation algorithms using Baker's approach

In this section we present in-place Logspace Approximation Scheme or LSAS ($(1 \pm \epsilon)$ approximation algorithm running in logspace for every $\epsilon > 0$) for a class of MSO-definable optimization problems which are amenable to Baker's method [12], also known as *shifting technique*, in locally bounded treewidth graphs.

There are two main general approaches for designing PTASs for problems on planar graphs. The first approach is based on planar separators [68]. The approximation algorithms resulting from this approach are impractical in general. To address this, Baker [12] introduced the second approach for PTASs in planar graphs, based on decomposition into overlapping subgraphs of bounded outerplanarity, which are of bounded treewidth. For a general account on these, see [45].

Baker's method was originally designed to give PTASs for a host of NP-hard optimization problems on planar graphs like minimum vertex cover, minimum dominating set, maximum independent set etc which are hard to approximate in general graphs. Many of these remain NP-hard even in planar graphs. Later the technique was generalized to a broader class of graphs called graphs of bounded local treewidth [54,44]. For a vertex v of a graph G and integer $k \geq 0$, by G_v^k we denote the subgraph of G induced by vertices within distance at most k from v in G . A class of graphs \mathcal{G} is of bounded local treewidth if there exists function f such that for every graph $G \in \mathcal{G}$ and every vertex v of G , $\text{treewidth}(G_v^k) \leq f(k)$.

8.2.1. Baker's algorithm

The main two computational bottlenecks in Baker's approach are

1. decomposing the graph into bounded treewidth graphs, and
2. solving the optimization problem on bounded tree width graphs optimally and combining these solutions.

Step (1) requires performing BFS on the graph G and considering the induced subgraphs $G_{i,j}$ between layers $ki + j$ and $k(i + 1) + j$ (which are of treewidth $O(k)$) for $i \geq 0$ and offset $0 \leq j \leq k - 1$ by deleting (or including in both the adjacent slices) the vertices/edges in every $k = O(1/\epsilon)$ -th BFS layer (particular details differ depending on the problem). By choosing the right offset, we can make sure this affects the optimum solution at most by ϵ factor. For Step (2), given an MSO-definable problem, using Boadlander's [18] and Courcelle's theorem [39] on bounded treewidth graphs we can get the optimal solution. Baker's approach is highly efficient in terms of time complexity as this gives linear time approximation schemes but the scenario changes when considering the efficiency of the approach in terms of the space needed.

Though we can use the result of [48] for the second part of the algorithm which proves logspace analogue of Boadlander's and Courcelle's theorem, in the first part we need to compute distance. BFS can be performed in NL in general graphs and in $UL \cap \text{co-UL}$ for planar graphs [81]. Since $UL \subseteq NL$ is not known (or even believed) to be inside L, the NL bound ($UL \cap \text{co-UL}$ for planar graphs) for Baker's algorithm is the best known in ROM [41]. Recently [42] has given an LSAS for maximum matching in planar graphs and some more sparse graph classes (a problem amenable to Baker's method) but a logspace analogue of Baker's algorithm in general is not yet known.

Since in our in-place models (both rotate and implicit) we can overcome this hurdle by computing distance in L, we obtain the following result. Notice that though our in-place algorithms change the ordering of the vertices in an adjacency list (or array) the graph remains the same and so the distances between vertices.

Theorem 12. *In both the rotate and the implicit model MSO-definable optimization problems which are amenable to Baker's method has an LSAS in locally bounded treewidth graphs.*

8.3. Solving NP-hard problems in in-place models

We show how to solve some NP-hard graph problems using logspace and exponential time in the rotate and implicit models. In particular, this implies that problems such as vertex cover and dominating set can be solved in exponential time

using $O(\lg n)$ bits in both the models. In contrast, note that, no NP-hard problem can be solved in the ROM model using $O(\lg n)$ bits unless $P=NP$.

Similar to Fomin et al. [57], we define a class of graph problems in NP which we call *graph subset problems* where the goal is to find a subset of vertices satisfying some property. We show a meta theorem showing that a restricted class of graph subset problems that are in NP admit log-space exponential algorithms in the rotate and implicit models.

Given a graph G with its adjacency list, we encode a subset of the vertices as follows. For every vertex in the subset, we bring in the minimum labelled vertex among its neighbors to the front of the list, and for others, we keep a vertex with a higher label (than the minimum) at the front of the list. So it takes a linear time to check whether a vertex is in the subset. The algorithm enumerates all subsets (by this encoding) and simply verifies using the NP algorithm whether that subset satisfies the required property until it finds a subset satisfying the property or it has enumerated all the subsets. By a standard theorem in complexity theory [7], every problem in NP is actually verifiable by a log-space ROM and hence the overall space taken by our algorithm is only logarithmic. Note that our algorithm requires that the adjacency list of any vertex has at least two values, i.e. that the degree of any vertex is at least two. Thus we have

Theorem 13. *Any graph subset problem in NP can be solved using $O(\lg n)$ bits of extra space (and exponential time) in the rotate and implicit models in graphs G having minimum degree 2.*

Remark 1. We remark that the above idea can work for other graph problems that are not necessarily subset problems. For example, for testing hamiltonicity, we can simply explore all neighbors of a vertex (starting at the smallest labelled neighbor so we know when we have explored them all) in a systematic fashion simply encoding them into the adjacency list by moving the current neighbor to the front of the list, and test whether together they form a cycle of length n .

If the graphs have larger minimum degree (say at least $2\lg n$), we can even encode pointers in each adjacency list and using that we can even test for graph isomorphism in logarithmic space by explicitly maintaining the vertex mapping by these encoding pointers.

Remark 2. The minimum degree 2 restriction in the above theorem is not a serious restriction as for many problems (like vertex cover, dominating set and traveling salesperson problem), there are preprocessing routines that can handle (not necessarily in our model) and eliminate degree 1 vertices.

9. Concluding remarks

Our initial motivation was to get around the limitations of ROM to obtain a reasonable model for graphs in which we can obtain space efficient algorithms. We achieved that by introducing two new frameworks and obtained efficient (of the order of $O(n^3 \lg n)$) algorithms using $O(\lg n)$ bits of space for fundamental graph search procedures. Note that there could be other ways to define such in-place models, and it would be interesting to find such models with competitive algorithms for the problems we have studied in this paper. We also discussed various applications of our DFS/BFS results, and it is not surprising that many simple corollaries would follow as DFS/BFS being the backbone of so many graph algorithms. We showed that some of these results also translate to improved space efficient algorithms in ROM (by simulating the rotate model algorithms in ROM with one pointer per list). With some effort, we can obtain a log-space algorithm for minimum spanning tree. These results can be contrasted with the state of the art results in ROM that take almost linear bits for some of these problems other than having large runtime bounds. All our algorithms are conceptually simple, and as they don't use any heavy data structures, we believe that they are also practical to implement. Still, there are plenty of algorithmic graph problems to be studied in these models. We believe that our work is the first step towards this and will inspire further investigation into designing in-place algorithms for other graph problems. One future direction would be to improve the running time of our algorithms to make them more practical. Another natural direction would be to consider the case when both of our models are augmented with extra $O(s)$ bits of space for small enough s . Can we then design more graph algorithms space efficiently in this model? Can we also obtain time-space tradeoffs in this setting for natural algorithmic graph problems?

Surprisingly we could design log-space algorithm for some P-complete problems, and so it is important to understand the power of our models. Towards that we discovered that we can even obtain log-space algorithms for some NP-hard graph problems. More specifically, we defined *graph subset problems* and obtained log-space exponential time algorithms for problems belonging to this class. One interesting future direction would be to determine the exact computational power of these models along with exploring the horizon of interesting complexity theoretic consequences of problems in these models.

Finally, another interesting and challenging open question would be to devise efficient in-place DFS or BFS algorithm while assuming information-theoretically optimal or succinct representation of the input graph [1,2,26,55,74].

CRedit authorship contribution statement

All the authors contributed roughly equally in obtaining the technical results, as well as the writing of the paper.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] H. Acan, S. Chakraborty, S. Jo, K. Nakashima, K. Sadakane, S.R. Satti, Succinct representations of intersection graphs on a circle, in: 31st DCC, IEEE, 2021, pp. 123–132.
- [2] H. Acan, S. Chakraborty, S. Jo, S.R. Satti, Succinct data structures for families of interval graphs, in: 16th WADS, in: Lecture Notes in Computer Science, vol. 11646, Springer, 2019, pp. 1–13.
- [3] A. Aggarwal, R.J. Anderson, A random NC algorithm for depth first search, *Combinatorica* 8 (1) (1988) 1–12.
- [4] A. Aggarwal, R.J. Anderson, M. Kao, Parallel depth-first search in general directed graphs, *SIAM J. Comput.* 19 (2) (1990) 397–409.
- [5] N. Alon, Y. Matias, M. Szegedy, The space complexity of approximating the frequency moments, *J. Comput. Syst. Sci.* 58 (1) (1999) 137–147.
- [6] R.J. Anderson, E.W. Mayr, Parallelism and the maximal path problem, *Inf. Process. Lett.* 24 (2) (1987) 121–126.
- [7] S. Arora, B. Barak, *Computational Complexity - a Modern Approach*, Cambridge University Press, 2009.
- [8] T. Asano, K. Buchin, M. Buchin, M. Korman, W. Mulzer, G. Rote, A. Schulz, Reprint of: memory-constrained algorithms for simple polygons, *Comput. Geom.* 47 (3) (2014) 469–479.
- [9] T. Asano, T. Izumi, M. Kiyomi, M. Konagaya, H. Ono, Y. Otachi, P. Schweitzer, J. Tarui, R. Uehara, Depth-first search using $O(n)$ bits, in: 25th ISAAC, 2014, pp. 553–564.
- [10] T. Asano, D.G. Kirkpatrick, K. Nakagawa, O. Watanabe, $\tilde{O}(\sqrt{n})$ -space and polynomial-time algorithm for planar directed graph reachability, in: 39th MFCS, in: LNCS, vol. 8634, 2014, pp. 45–56.
- [11] T. Asano, W. Mulzer, G. Rote, Y. Wang, Constant-work-space algorithms for geometric problems, *JoCG* 2 (1) (2011) 46–68.
- [12] B.S. Baker, Approximation algorithms for np-complete problems on planar graphs, *J. ACM* 41 (1) (1994) 153–180.
- [13] N. Banerjee, S. Chakraborty, V. Raman, S. Roy, S. Saurabh, Time-space tradeoffs for dynamic programming in trees and bounded treewidth graphs, in: 21st COCOON, in: LNCS, vol. 9198, Springer, 2015, pp. 349–360.
- [14] N. Banerjee, S. Chakraborty, V. Raman, S.R. Satti, Space efficient linear time algorithms for BFS, DFS and applications, in: Theory of Computing Systems, Jan 2018.
- [15] L. Barba, M. Korman, S. Langerman, K. Sadakane, R.I. Silveira, Space-time trade-offs for stack-based algorithms, *Algorithmica* 72 (4) (2015) 1097–1129.
- [16] G. Barnes, J. Buss, W. Ruzzo, B. Schieber, A sublinear space, polynomial time algorithm for directed s-t connectivity, *SIAM J. Comput.* 27 (5) (1998) 1273–1282.
- [17] P. Beame, A general sequential time-space tradeoff for finding unique elements, *SIAM J. Comput.* 20 (2) (1991) 270–277.
- [18] H.L. Bodlaender, A linear-time algorithm for finding tree-decompositions of small treewidth, *SIAM J. Comput.* 25 (6) (1996) 1305–1317.
- [19] A. Borodin, S.A. Cook, A time-space tradeoff for sorting on a general sequential model of computation, *SIAM J. Comput.* 11 (2) (1982) 287–297.
- [20] A. Borodin, M.J. Fischer, D.G. Kirkpatrick, N.A. Lynch, M. Tompa, A time-space tradeoff for sorting on non-oblivious machines, *J. Comput. Syst. Sci.* 22 (3) (1981) 351–364.
- [21] H. Brönnimann, T.M. Chan, E.Y. Chen, Towards in-place geometric algorithms and data structures, in: 20th SOCG, 2004, pp. 239–246.
- [22] H. Bührman, R. Cleve, M. Koucký, B. Loff, F. Speelman, Computing with a full memory: catalytic space, in: STOC, 2014, pp. 857–866.
- [23] H. Bührman, M. Koucký, B. Loff, F. Speelman, Catalytic space: non-determinism and hierarchy, in: 33rd STACS, 2016, pp. 24:1–24:13.
- [24] D. Chakraborty, A. Pavan, R. Tewari, N.V. Vinodchandran, L. Yang, New time-space upperbounds for directed reachability in high-genus and h-minor-free graphs, in: FSTTCS, 2014, pp. 585–595.
- [25] S. Chakraborty, Space Efficient Graph Algorithms, PhD thesis, The Institute of Mathematical Sciences, HBNI, India, 2018.
- [26] S. Chakraborty, S. Jo, K. Sadakane, S.R. Satti, Succinct data structures for small clique-width graphs, in: 31st DCC, IEEE, 2021, pp. 133–142.
- [27] S. Chakraborty, S. Jo, S.R. Satti, Improved space-efficient linear time algorithms for some classical graph problems, CoRR, arXiv:1712.03349 [abs], 2017.
- [28] S. Chakraborty, A. Mukherjee, V. Raman, S.R. Satti, A framework for in-place graph algorithms, in: 26th ESA, 2018, pp. 13:1–13:16.
- [29] S. Chakraborty, A. Mukherjee, S.R. Satti, Space efficient algorithms for breadth-depth search, in: 22nd FCT, 2019, pp. 201–212.
- [30] S. Chakraborty, V. Raman, S.R. Satti, Biconnectivity, st-numbering and other applications of DFS using $O(n)$ bits, *J. Comput. Syst. Sci.* 90 (2017) 63–79.
- [31] S. Chakraborty, K. Sadakane, Indexing graph search trees and applications, in: 44th MFCS, 2019, pp. 67:1–67:14.
- [32] S. Chakraborty, K. Sadakane, S.R. Satti, Optimal in-place algorithms for basic graph problems, in: 31st IWOC, in: Lecture Notes in Computer Science, vol. 12126, Springer, 2020, pp. 126–139.
- [33] S. Chakraborty, S.R. Satti, Space-efficient algorithms for maximum cardinality search, its applications, and variants of BFS, *J. Comb. Optim.* (Mar 2018).
- [34] T.M. Chan, E.Y. Chen, Multi-pass geometric algorithms, *Discrete Comput. Geom.* 37 (1) (2007) 79–102.
- [35] T.M. Chan, J.I. Munro, V. Raman, Faster, space-efficient selection algorithms in read-only memory for integers, in: 24th ISAAC, 2013, pp. 405–412.
- [36] T.M. Chan, J.I. Munro, V. Raman, Selection and sorting in the “restore” model, in: 25th-SODA, 2014, pp. 995–1004.
- [37] S.A. Cook, C. Rackoff, Space lower bounds for maze threadability on restricted machines, *SIAM J. Comput.* 9 (3) (1980) 636–652.
- [38] T.H. Cormen, C.E. Leiserson, R.L. Rivest, C. Stein, *Introduction to Algorithms*, 3. ed., MIT Press, 2009.
- [39] B. Courcelle, M. Mosbah, Monadic second-order evaluations on tree-decomposable graphs, *Theor. Comput. Sci.* 109 (1–2) (1993) 49–82.
- [40] O. Darwish, A. Elmasry, Optimal time-space tradeoff for the 2d convex-hull problem, in: 22th ESA, 2014, pp. 284–295.
- [41] S. Datta, R. Kulkarni, Space complexity of optimization problems in planar graphs, in: 11th TAMC, 2014, pp. 300–311.
- [42] S. Datta, R. Kulkarni, A. Mukherjee, Space-efficient approximation scheme for maximum matching in sparse graphs, in: 41st MFCS, 2016, pp. 28:1–28:12.
- [43] S. Datta, N. Limaye, P. Nimbhorkar, T. Thierauf, F. Wagner, Planar graph isomorphism is in log-space, in: 24th CCC, 2009, pp. 203–214.
- [44] E.D. Demaine, M. Taghi Hajiaghayi, Equivalence of local treewidth and linear local treewidth and its algorithmic applications, in: 15th SODA, 2004, pp. 840–849.
- [45] E.D. Demaine, M. Taghi Hajiaghayi, *Approximation Schemes for Planar Graph Problems*, Springer US, Boston, MA, 2008, pp. 1–99.
- [46] Y. Dodis, M. Patrascu, M. Thorup, Changing base without losing space, in: 42nd STOC, 2010, pp. 593–602.
- [47] J. Edmonds, C.K. Poon, D. Achlioptas, Tight lower bounds for st-connectivity on the NNJAG model, *SIAM J. Comput.* 28 (6) (1999) 2257–2284.
- [48] M. Elberfeld, A. Jakob, T. Tantau, Logspace versions of the theorems of bodlaender and courcelle, in: 51th FOCS, 2010, pp. 143–152.
- [49] M. Elberfeld, K. Kawarabayashi, Embedding and canonizing graphs of bounded genus in logspace, in: STOC, 2014, pp. 383–392.
- [50] M. Elberfeld, P. Schweitzer, Canonizing graphs of bounded tree width in logspace, in: 33rd STACS, 2016, pp. 32:1–32:14.
- [51] A. Elmasry, Why depth-first search efficiently identifies two and three-connected graphs, in: 21st ISAAC, 2010, pp. 375–386.
- [52] A. Elmasry, T. Hagerup, F. Kammer, Space-efficient basic graph algorithms, in: 32nd STACS, 2015, pp. 288–301.
- [53] A. Elmasry, D.D. Juhl, J. Katajainen, S.R. Satti, Selection from read-only memory with limited workspace, *Theor. Comput. Sci.* 554 (2014) 64–73.

- [54] D. Eppstein, Diameter and treewidth in minor-closed graph families, *Algorithmica* 27 (3) (2000) 275–291.
- [55] A. Farzan, J.I. Munro, Succinct encoding of arbitrary graphs, *Theor. Comput. Sci.* 513 (2013) 38–52.
- [56] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, J. Zhang, On graph problems in a semi-streaming model, *Theor. Comput. Sci.* 348 (2–3) (2005) 207–216.
- [57] F.V. Fomin, S. Gaspers, D. Lokshantov, S. Saurabh, Exact algorithms via monotone local search, in: *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18–21, 2016*, 2016, pp. 764–775.
- [58] G. Franceschini, J.I. Munro, Implicit dictionaries with $O(1)$ modifications per update and fast search, in: *17th SODA, 2006*, pp. 404–413.
- [59] G. Franceschini, S. Muthukrishnan, In-place suffix sorting, in: *34th ICALP, 2007*, pp. 533–545.
- [60] G. Franceschini, S. Muthukrishnan, M. Patrascu, Radix sorting with no extra space, in: *Algorithms – ESA 2007, 15th Annual European Symposium, Eilat, Israel, October 8–10, 2007, Proceedings, 2007*, pp. 194–205.
- [61] G.N. Frederickson, Upper bounds for time-space trade-offs in sorting and selection, *J. Comput. Syst. Sci.* 34 (1) (1987) 19–26.
- [62] T. Hagerup, F. Kammer, Succinct choice dictionaries, *CoRR*, arXiv:1604.06058 [abs], 2016.
- [63] F. Kammer, D. Kratsch, M. Laudahn, Space-efficient biconnected components and recognition of outerplanar graphs, in: *41st MFCS, 2016*.
- [64] F. Kammer, A. Sajenko, Linear-time in-place DFS and BFS in the restore model, *CoRR*, arXiv:1803.04282 [abs], 2018.
- [65] Donald E. Knuth, *The Art of Computer Programming, Volume I: Fundamental Algorithms*, 2nd edition, Addison-Wesley, 1973.
- [66] M. Koucký, Catalytic computation, *Bull. Eur. Assoc. Theor. Comput. Sci.* 118 (2016).
- [67] T.W. Lai, D. Wood, Implicit selection, in: *1st SWAT, 1988*, pp. 14–23.
- [68] R.J. Lipton, R.E. Tarjan, Applications of a planar separator theorem, *SIAM J. Comput.* 9 (3) (1980) 615–627.
- [69] A. McGregor, Graph stream algorithms: a survey, *SIGMOD Rec.* 43 (1) (2014) 9–20.
- [70] J.I. Munro, An implicit data structure supporting insertion, deletion, and search in $o(\log^2 n)$ time, *J. Comput. Syst. Sci.* 33 (1) (1986) 66–74.
- [71] J.I. Munro, Tables, in: *FSTTCS, 1996*, pp. 37–42.
- [72] J.I. Munro, M. Paterson, Selection and sorting with limited storage, *Theor. Comput. Sci.* 12 (1980) 315–323.
- [73] J.I. Munro, V. Raman, Selection from read-only memory and sorting with minimum data movement, *Theor. Comput. Sci.* 165 (2) (1996) 311–323.
- [74] J.I. Munro, K. Wu, Succinct data structures for chordal graphs, in: *29th ISAAC, in: LIPIcs, vol. 123, Schloss Dagstuhl – Leibniz-Zentrum Für Informatik, 2018*, pp. 67:1–67:12.
- [75] J. Pagter, T. Rauhe, Optimal time-space trade-offs for sorting, in: *39th FOCS, 1998*, pp. 264–268.
- [76] J.H. Reif, Symmetric complementation, *J. ACM* 31 (2) (1984) 401–421.
- [77] J.H. Reif, Depth-first search is inherently sequential, *Inf. Process. Lett.* 20 (5) (1985) 229–234.
- [78] O. Reingold, Undirected connectivity in log-space, *J. ACM* 55 (4) (2008).
- [79] T. Tantau, Logspace optimization problems and their approximability properties, *Theory Comput. Syst.* 41 (2) (2007) 327–350.
- [80] R.E. Tarjan, Space-efficient implementations of graph search methods, *ACM Trans. Math. Softw.* 9 (3) (1983) 326–339.
- [81] T. Thierauf, F. Wagner, The isomorphism problem for planar 3-connected graphs is in unambiguous logspace, *Theory Comput. Syst.* 47 (3) (2010) 655–673.
- [82] M. Tompa, Two familiar transitive closure algorithms which admit no polynomial time, sublinear space implementations, *SIAM J. Comput.* 11 (1) (1982) 130–137.