

# Sampling Arborescences in Parallel

**Nima Anari**

Stanford University, CA, USA  
anari@cs.stanford.edu

**Nathan Hu**

Stanford University, CA, USA  
zixia314@stanford.edu

**Amin Saberi**

Stanford University, CA, USA  
saber@stanford.edu

**Aaron Schild**

University of Washington, Seattle, WA, USA  
aschild@berkeley.edu

---

## Abstract

We study the problem of sampling a uniformly random directed rooted spanning tree, also known as an arborescence, from a possibly weighted directed graph. Classically, this problem has long been known to be polynomial-time solvable; the exact number of arborescences can be computed by a determinant [33], and sampling can be reduced to counting [18, 16]. However, the classic reduction from sampling to counting seems to be inherently sequential. This raises the question of designing efficient parallel algorithms for sampling. We show that sampling arborescences can be done in RNC.

For several well-studied combinatorial structures, counting can be reduced to the computation of a determinant, which is known to be in NC [9]. These include arborescences, planar graph perfect matchings, Eulerian tours in digraphs, and determinantal point processes. However, not much is known about efficient parallel sampling of these structures. Our work is a step towards resolving this mystery.

**2012 ACM Subject Classification** Theory of computation → Parallel algorithms; Theory of computation → Generating random combinatorial structures; Theory of computation → Random walks and Markov chains

**Keywords and phrases** parallel algorithms, arborescences, spanning trees, random sampling

**Digital Object Identifier** 10.4230/LIPIcs.ITCS.2021.83

**Related Version** *Full Version:* <https://arxiv.org/abs/2012.09502>.

## 1 Introduction

Algorithms for (approximately) counting various combinatorial structures are often based on the equivalence between (approximate) counting and sampling [18, 16]. This is indeed the basis of the Markov Chain Monte Carlo (MCMC) method to approximate counting, which is arguably the most successful approach to counting, resolving long-standing problems such as approximating the permanent [17] and computing the volume of convex sets [13].

Approximate sampling and counting are known to be equivalent for a wide class of problems, including the so-called self-reducible ones [18, 16]. This equivalence is nontrivial and most useful in the direction of reducing counting to sampling. However, for some problems, the “easier” direction of this equivalence, namely the reduction from sampling to counting, proves useful. For these problems, almost by definition, we can count via approaches other than MCMC.



© Nima Anari, Nathan Hu, Amin Saberi, and Aaron Schild;  
licensed under Creative Commons License CC-BY

12th Innovations in Theoretical Computer Science Conference (ITCS 2021).

Editor: James R. Lee; Article No. 83; pp. 83:1–83:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

One of the mysterious approaches to counting is via determinant computations. A range of counting problems can be (exactly) solved by simply computing a determinant. A non-exhaustive list is provided below.

- Spanning trees in a graph can be counted by computing a determinant related to the Laplacian of the graph, a result known as the matrix-tree theorem [22].
- Arborescences in a directed graph can be counted by computing a determinant related to the directed Laplacian [33].
- The number of perfect matchings in a planar graph can be computed as the Pfaffian (square root of the determinant) of an appropriately signed version of the adjacency matrix, a.k.a. the Tutte matrix [20].
- The number of Eulerian tours in an Eulerian digraph is directly connected to the number of arborescences, and consequently the determinant related to the directed Laplacian [1, 34].
- Given vectors  $v_1, \dots, v_n \in \mathbb{R}^d$ , the volume sampling distribution on subsets  $S \in \binom{[n]}{d}$  can be defined as follows:

$$\mathbb{P}[S] \propto \det([v_i]_{i \in S})^2.$$

The partition function of this distribution is simply  $\det(\sum_i v_i v_i^T)$  (see, e.g., [10]).

- The number of non-intersecting paths between specified terminals in a lattice, and more generally applications of the Lindström-Gessel-Viennot lemma [24, 14].

Efficient counting for these problems follows the polynomial-time computability of the associated determinants. In turn, one obtains efficient sampling algorithms for all of these problems; we remark that not all of these problems are known to be self-reducible, but nevertheless “easy” slightly varied sampling to counting reductions exist for all of them.

While polynomial-time sampling for all of these problems has long been settled, we reopen the investigation of these problems by considering *efficient parallel sampling* algorithms. We focus on the computational model of PRAM, and specifically on the complexity classes NC and RNC. Here, a polynomially bounded number of processors are allowed access to a shared memory, and the goal is for the running time to be polylogarithmically bounded; the class RNC has additionally access to random bits. Determinants can be computed efficiently in parallel, in the class NC [9], and as a result there are NC counting algorithms for all of aforementioned problems. However, the sampling to counting reductions completely break down for parallel algorithms, as there seems to be an inherent sequentiality in these reductions.

Take spanning trees in a graph as an example. The classic reduction from sampling to counting proceeds as follows:

```

for each edge  $e$  do
   $A \leftarrow$  number of spanning trees containing  $e$ 
   $B \leftarrow$  total number spanning trees
  Flipping a coin with bias  $A/B$ , decide whether  $e$  should be part of the tree.
  Either contract or delete the edge  $e$  based on this decision.

```

Each iteration of this loop uses a counting oracle to compute  $A, B$ . However the decision of whether to include an edge  $e$  as part of the tree affects future values of  $A, B$  for other edges, and this seems to be the inherent sequentiality in this algorithm. The sampling-to-counting reduction for all other listed problems encounters the same sequentiality obstacle.

In this paper, we take a step towards resolving the mysterious disparity between counting and sampling in the parallel algorithms world. We resolve the question of sampling arborescences in weighted directed graphs, and as a special case spanning trees in weighted undirected graphs, using efficient parallel algorithms with access to randomness, a.k.a. the class RNC.

We remark that the special case of sampling spanning trees in unweighted undirected graphs was implicitly solved by the work of Teng [30], who showed how to simulate random walks in RNC. When combined with earlier work of Aldous [2] and Broder [6], this algorithm would simulate a random walk on the graph, and from its transcript extract a random spanning tree. However, adding either weights or directions to the graph results in the need for potentially exponentially large random walks, which cannot be done in RNC. Our work removes this obstacle.

► **Theorem 1.** *There is an RNC algorithm which takes a directed graph  $G = (V, E)$  together with edge weights  $w : E \rightarrow \mathbb{R}_{\geq 0}$  as input and outputs a random directed rooted tree  $T$ , a.k.a. an arborescence. The output  $T$  follows the distribution*


$$\mathbb{P}[T] \propto \prod_{e \in T} w(e).$$

## 1.1 Related Work and Techniques


There is a long line of research on algorithms for sampling and counting spanning trees and more generally arborescences. The matrix-tree theorem of Kirchhof [22] showed how to count spanning trees in undirected graphs, and later Tutte [33] generalized this to arborescences in digraphs. Somewhat surprisingly Aldous [2] and Broder [6] showed that random spanning trees and more generally random arborescences of a graph can be extracted from the transcript of a random walk on the graph itself. The main focus of subsequent work on this problem has been on improving the total running time of sequential algorithms for sampling. After a long line of work [35, 7, 21, 25, 12, 11], Schild [28] obtained the first almost-linear time algorithm for sampling spanning trees. More recently Anari, Liu, Oveis Gharan, and Vinzant [3] improved this to nearly-linear time. Many of these works are based on speeding up the Aldous-Broder algorithm. Our main result, Theorem 1, is also built on the Aldous-Broder algorithm, but we focus on parallelizing it instead of optimizing the total running time.

No almost-linear time algorithm is yet known for sampling arborescences in digraphs, as opposed to spanning trees in undirected graphs. While sampling spanning trees has a multitude of application (see [28]), there are a number of applications for the directed graph generalization. Most notably, there is a many-to-one direct correspondence between Eulerian tours in an Eulerian digraph and arborescences of the graph. This correspondence, known as the BEST theorem [1, 34] allows one to generate random Eulerian tours by generating random arborescences (see [8]). We leave the question of whether the correspondence in the BEST theorem is implementable in NC to future work, but note that sampling Eulerian tours has interesting applications in biology and sequence processing [19, 27]. We remark that, unlike directed graphs, generating random Eulerian tours of *undirected* Eulerian graphs in polynomial time is a major open problem [32].

In a slightly different direction related to this work, Balaji and Datta [4] considered the space complexity of counting arborescences. They showed this problem is in L for graphs of bounded tree-width, obtaining algorithms for counting Eulerian tours in these digraphs as well.



■ **Figure 1** Starting from left it takes  $\Theta(W)$  steps to cover.



■ **Figure 2** A random walk started from the left node covers all  $n$  nodes in time  $\Theta(2^n)$ .

Perhaps in the first major result of its kind in the search for RNC sampling algorithms, Teng [31] showed implicitly how to sample spanning trees in undirected graphs in RNC. Teng [31] showed how to parallelize the simulation of a random walk on a graph; more precisely, he showed how to output a length  $L$  trace of a walk on size  $n$  Markov chains in parallel running time  $\text{polylog}(L, n)$  using only  $\text{poly}(L, n)$  many processors. The Aldous-Broder algorithm extracts an arborescence from the trace of a random walk by extracting the so-called first-visit edge to each vertex. If a random walk is simulated until all vertices are visited at least once, the trace of the random walk has enough information to extract all such first-visit edges. This allows RNC sampling of arborescences in graphs where the number of steps needed to visit all vertices, known as the cover time, is polynomially bounded. While the cover time of a random walk on an undirected unweighted graph is polynomially bounded in the size of the graph (see, e.g., [23]), adding either weights or directions can make the cover time exponentially large. See Figures 1 and 2.

We overcome the obstacle of exponentially large cover times, by taking a page from some of the recent advances on the sequential sampling algorithms for spanning trees. Instead of simulating the entirety of the random walk until cover time, we extract only the first-visit edges to each vertex. We use the same insight used in several prior works that once a region of the graph has been visited, subsequent visits of the random walk can be *shortcut* to the first edge that exits this region. However, this involves a careful construction of a hierarchy of “regions”, and bounding the number of steps needed inside each region to fully cover it. Unlike undirected graphs, in directed graphs arguing about the covering time of a region becomes complicated. We build on some of the techniques developed by Boczkowski, Peres, and Sousi [5] to bound these covering times in the case of Eulerian digraphs. We then reduce the arborescence sampling problem for arbitrary digraphs to that of Eulerian digraphs.

## 1.2 Overview of the Algorithm

At a high-level, our algorithm first proceeds by reducing the problem to sampling an arborescence from an Eulerian graph. We then construct a “loose decomposition” of the Eulerian graph into weighted cycles. We then construct a hierarchy of vertex sets by starting with the empty graph and adding the weighted cycles one by one, from the lowest weight to the highest, adding the connected components in each iteration to the hierarchy. This results in a hierarchical clustering of vertices, with the intuitive property that once we enter a cluster, we spend a lot of time exploring inside before exiting. Our algorithm proceeds by “simulating” polynomially many jumps between the children of each cluster, before exiting that cluster. The resulting jumps are then stitched together at all levels of the hierarchy to form a partial subset of the transcript of a random walk.

The shallow “simulation” of edges jumping across children of a cluster in the hierarchy can be done in RNC by using a doubling trick as in the work of Teng [31]; in order to “simulate”  $L$  jumping edges, we combine the first  $L/2$  with the last  $L/2$  in a recursive fashion. A naive

implementation of this is not parallelizable however, as one needs to know *where* the first  $L/2$  jumps eventually land before “simulating” the rest. However since the number of locations is polynomially bounded, one can precompute an answer for each possible landing location in parallel to simulating the first  $L/2$  edges. Once shallow jumps are simulated, our algorithm then proceeds by stitching these jumping edges with jumping edges of one level lower (i.e., inside of the children clusters), and so on. Again, a naive implementation of this stitching is not in RNC, but we employ a similar doubling trick and an additional caching trick to parallelize everything. We then argue that with high probability all of the edges extracted by this algorithm contain the first-visit edges to every vertex.

## 2 Preliminaries

We use the notation  $[n]$  to denote the set  $\{1, \dots, n\}$ . We use the notation  $\tilde{O}(\cdot)$  to hide polylogarithmic factors. When we use the term high probability, we mean with probability  $1 - \text{poly}(\frac{1}{n})$ , where  $n$  is the size of the input, and the polynomial can be taken arbitrarily large by appropriately setting parameters.

### 2.1 Graph Theory Notations

When a graph  $G = (V, E)$  is clear from the context, we use  $n$  to refer to  $|V|$  and  $m$  to refer to  $|E|$ .

For a subset of vertices  $S \subset V$ ,  $\delta^+(S)$  denotes the set all incoming edges  $\{e = (u, v) \mid v \in S, u \notin S\}$ . Similarly,  $\delta^-(S)$  will denote the set of all outgoing edges  $\{e = (u, v) \mid u \in S, v \notin S\}$ . Lastly, let  $G(S)$  denote the subgraph induced by  $S$ .

► **Definition 2.** *Given a weighed graph  $G = (V, E)$  with weights  $w : E \rightarrow \mathbb{R}_{\geq 0}$ , a subset of the vertices  $S \subseteq V$  is said to be strongly connected by edges of weight  $w_c$  if for every  $s, t, \in S$ , there exists a path of edges inside  $S$  connecting  $s$  to  $t$  such that every edge  $e$  of the path has weight  $w(e) \geq w_c$ .*

► **Definition 3.** *An arborescence is a directed graph rooted at vertex  $r$  such that for any other vertex  $v$ , there is exactly one path from  $v$  to  $r$ .*

Arborescences are also known as directed rooted trees and are a natural analog of spanning trees in directed graphs. When a background graph  $G = (V, E)$  is clear from context, by an arborescence we mean a subgraph of  $G$  that is an arborescence.

Given digraph  $G = (V, E)$  together with a weight function  $w : E \rightarrow \mathbb{R}_{> 0}$ , the random arborescence distribution is the distribution that assigns to each arborescence  $T$  probability

$$\mathbb{P}[T] \propto \prod_{e \in T} w(e).$$

When  $w$  is not given, it is assumed to be the constant 1 function, and the random arborescence distribution becomes uniformly distributed over all arborescences. We view an undirected graph  $G = (V, E)$  as a directed graph with double the number of edges, with each undirected edge producing two directed copies in the two possible directions. It is easy to see that the weighted/uniform random arborescence distribution on the directed version of an undirected graph is the same as the weighted/uniform spanning tree distribution on the original graph; viewing each arborescence without directions yields the corresponding spanning tree.

We call a (possibly weighted) digraph  $G = (V, E)$  Eulerian if

$$\sum_{e \in \delta^+(v)} w(e) = \sum_{e \in \delta^-(v)} w(e)$$

for all vertices  $v$ . Note that the directed version of an undirected graph is automatically Eulerian.

## 2.2 Random Walks and Arborescences

Our approach centers around the Aldous-Broder algorithm on directed graphs [2, 6]. Their work reduces the task of randomly generating a spanning tree or arborescence to simulating a random walk on the graph until all vertices have been visited. While most famously known for generating spanning trees, this algorithm can be used to generate arborescences as well.

A weighted digraph defines a natural random walk. This is a Markov process  $X_0, X_1, \dots$ , where each  $X_i$  is obtained as a random neighbor of  $X_{i-1}$  by choosing an outgoing edge with probability proportional to its weight, and transitioning to its endpoint.

$$\mathbb{P}[X_i \mid X_{i-1}] \propto w(X_{i-1}, X_i).$$

A stationary distribution  $\pi$  is a distribution on the vertices such that if  $X_0$  is chosen according to  $\pi$ , then all  $X_i$  are distributed as  $\pi$ . Under mild conditions, namely strong connectivity and aperiodicity, the stationary distribution is unique.

A random walk on an undirected graph is known to be time-reversible. That is if  $X_0$  is started from the stationary distribution, then  $(X_i, X_{i+1})$  is identical in distribution to  $(X_{i+1}, X_i)$ . This does not hold for directed graphs. However the time reversal of the process  $\dots, X_i, X_{i+1}, \dots$  corresponds to a random walk on a different digraph.

► **Definition 4.** For a weighted digraph  $G = (V, E)$  with weights  $w : E \rightarrow \mathbb{R}_{\geq 0}$  and stationary distribution  $\pi$ , define the time-reversal to be the graph  $G' = (V, E')$  on the same set of vertices, with edges reversed in direction, and weights given by

$$w'(v, u) = \pi(u)w(u, v)/\pi(v).$$

The random walk on  $G'$  shares the same stationary distribution  $\pi$  as  $G$ . The random walk on  $G'$  (started from  $\pi$ ) is identical in distribution to the time-reversed random walk on  $G$  (started from  $\pi$ ); see [23].

► **Theorem 5** ([2, 6]). Suppose that  $G = (V, E)$  is a strongly connected weighted graph, with weights given by  $w : E \rightarrow \mathbb{R}_{> 0}$ . Perform a time-reversed random walk, starting from a vertex  $r$ , until all vertices are visited at least once. For each vertex  $v \in V \setminus \{r\}$ , record the edge used in the random walk to reach  $v$  for the first time. Let  $T$  be the collection of all these first-visit edges (with directions reversed). Then  $T$  is an arborescence of  $G$  rooted at  $r$ , and

$$\mathbb{P}[T] \propto \prod_{e \in T} w(e)$$

This allows us to sample  $r$ -rooted arborescence from a digraph by performing a random walk on the time-reversal. It is also known that among all arborescences, the total weight of those rooted at  $r$  is proportional to  $\pi(r)$ , where  $\pi$  is the stationary distribution [2, 6]. Therefore to resolve Theorem 1, it is enough to show how to sample  $r$ -rooted arborescences.

► **Theorem 6.** There is an RNC algorithm which takes a directed graph  $G = (V, E)$  together with edge weights  $w : E \rightarrow \mathbb{R}_{\geq 0}$  and a root vertex  $r \in V$  as input and outputs a random  $r$ -rooted arborescence, where

$$\mathbb{P}[T] \propto \prod_{e \in T} w(e).$$

We can fix the root because the stationary distribution  $\pi$  can be computed in NC, by solving a system of linear equations.

## 2.3 Parallel Algorithms

In this paper we consider parallel algorithms run on the PRAM model, and we will construct algorithms to show that sampling problems are in the class RNC, the randomized version of NC. In this class, we are allowed to use polynomially (in the input size) many processors who share access to a common random access memory, and also have access to random bits. The running time of our algorithms must be polylogarithmic in the input size.

While our work hinges upon the simulation of a random walk, a process which is inherently sequential, for polynomially many steps, such a task is known to be in RNC through the use of a “doubling trick.”

► **Theorem 7** ([31]). *Suppose that  $G = (V, E)$  is a directed graph, with weights given by  $w : E \rightarrow \mathbb{R}_{\geq 0}$ . The transcript of a random walk starting from a given vertex  $v_0 \in V$  and running for  $L$  time steps can be produced in the PRAM model using  $\text{poly}(|E|, L)$  processors in  $\text{polylog}(|E|, L)$  time.*

Theorem 5 and Theorem 7 naturally lead to a parallel algorithm that randomly samples arborescences using  $\text{poly}(|E|, L)$  processors and  $\text{polylog}(|E|, L)$  time where  $L$  is the cover time of the digraph. However,  $L$  can be exponential in the size of the graph as the cover time depends closely on the edge weights of a graph. In directed graphs, the cover time can be exponential even if we do not allow weights. See Figures 1 and 2.

Luckily, Theorem 5 only needs the first-visit edges to produce the random arborescence. This insight has been heavily used to improve the running time of sequential algorithms for sampling spanning trees [21, 25, 28], by *shortcutting* the random walk. Here we use the same insight to design an RNC algorithm. Our algorithm identifies a hierarchy of clusters of vertices  $S \subseteq V$  and will only simulate incoming and outgoing edges of each cluster as opposed to the entire walk. To do this, we will use a well-known primitive that the probability of entering or exiting a cluster through any edge can be computed using a system of linear equations.

► **Lemma 8** (see, e.g., [28]). *Given a set  $S \subset V$  and vertex  $v \in S$ , the probability of a random walk started at  $v$  exiting  $S$  through any particular edge  $e \in \delta^-(S)$  can be computed by solving a system of linear equations involving the Laplacian.*

Some of the most powerful primitives for NC algorithms come from linear algebra. In particular, multiplying matrices, computing determinants, and inverting matrices all have NC algorithms [9]. Combining this with Lemma 8 we obtain the following folklore result.

► **Lemma 9.** *Given a set  $S \subseteq V$  and vertex  $v \in S$ , the probability of a random walk started at  $v$  exiting  $S$  through any particular edge  $e \in \delta^-(S)$  can be computed in NC.*

## 2.4 Schur Complements

A key tool which we shall use in our analysis of random walks is Schur complements.

► **Definition 10.** *For any weighted digraph  $G = (V, E)$ , for any subset of the vertices  $S \subseteq V$ , we define the graph  $G_S$  to be the Schur complement of  $\bar{S}$ .  $G_S$  is the directed graph formed by starting with the induced sub-graph of  $S$ . Then, for every pair of vertices  $(u, v)$  in  $S$ , add an edge from  $u$  to  $v$  with weight  $\deg(u) \mathbb{P}[u \rightarrow v]$ , where  $\mathbb{P}[u \rightarrow v]$  is the probability that a random walk on  $G$  currently at  $u$  exits  $S$  with the next move and its first return to  $S$  is at the vertex  $v$ .*

Then, it is easy to see that for any starting vertex  $v_0 \in S$ , the distribution over random walk transcripts in  $G_S$  starting at  $v_0$  is the same as the distribution over random walk transcripts in  $G$  starting at  $v_0$  with all vertices outside of  $S$  removed. We also observe that:

► **Proposition 11.** *For any Eulerian weighted digraph  $G = (V, E)$ , for any subset of the vertices  $S \subset V$ , the Schur complement  $G_S$  is also an Eulerian weighted digraph and the degree of any vertex in  $S$  is the same in  $G_S$  as in  $G$ .*

### 3 Reduction to Eulerian Graphs

A key part of our algorithm is the analysis of the time it takes for random walks to cover regions of a digraph. Cover times are in general difficult to analyze on directed graphs, but the situation becomes much easier on Eulerian graphs. For more precise statements, see Section 6. As a first step in our algorithm, we show how to reduce the design of sampling algorithms on arbitrary digraphs to Eulerian ones.

We can reduce the problem of sampling an arborescence of digraph  $G$  to sampling a random arborescence rooted at vertex  $r$  on strongly connected Eulerian digraph  $G''$  as follows in Algorithm 1. We begin by selecting the vertex at which the arborescence is rooted, using the Markov chain tree theorem [2, 6].

► **Theorem 12** (The Markov chain tree theorem). *Let  $G = (V, E)$  be a weighted directed graph. Assume that the natural Markov chain associated with  $G$  has stationary distribution  $\pi = (\pi_1, \dots, \pi_n)$ . Then we have the following relationship between  $\pi$  and the arborescences of  $G$ :*

$$\pi_i = \frac{\sum \{w(T) \mid T \text{ arborescence rooted at } i\}}{\sum \{w(T) \mid T \text{ arborescence}\}},$$

where  $w(T)$  is product of edge weights in  $T$ .

Finding the stationary distribution of a random walk on some graph  $G$  reduces to solving a system of  $n$  linear equations and can be done in NC.

After the root  $r \in V$  of an arborescence is selected, we can add outgoing edges from  $r$  to graph  $G$  to guarantee that  $G$  is strongly connected. As no arborescence rooted at  $r$  will contain these new edges, the addition of such edges does not affect the distribution of arborescences rooted at  $r$ . Then, we note that for any vertex  $v$ , if we multiply the weights of all edges in  $\delta^-(v)$  by some constant, the distribution of arborescences rooted at any vertex does not change. We can therefore rescale edges to ensure that the resulting graph is Eulerian.

► **Proposition 13.** *Graph  $G''$  is Eulerian.*

**Proof.** We have that, for all  $v$  in  $G'$

$$\sum_{(u,v) \in \delta^+(v)} \pi'(u) \frac{w'(u,v)}{\deg'_{\text{out}}(u)} = \pi'(v)$$

Then, considering  $G''$  :

$$\begin{aligned} \deg''_{\text{in}}(v) &= \sum_{(u,v) \in \delta^+(v)} w''(u,v) = \sum_{(u,v) \in \delta^+(v)} \frac{w'(u,v)\pi'(u)}{\deg'_{\text{out}}(u)} = \pi'(v) \\ &= \pi'(v) \sum_{(v,x) \in \delta^-(v)} \frac{w'(v,x)}{\deg'_{\text{out}}(v)} = \sum_{(v,x) \in \delta^-(v)} w''(v,x) = \deg''_{\text{out}}(v) \quad \blacktriangleleft \end{aligned}$$



■ **Algorithm 1** Reduction to strongly-connected Eulerian graphs.

---

```

Compute the stationary distribution  $\pi(v)$  of a random walk on  $G$ 
Choose vertex  $r$  to be the root of the arborescence with probability  $\pi(r)$ 
 $G' = (V, E') \leftarrow G$ 
for  $v \in V \setminus \{r\}$  in parallel do
  | if  $e' = (r, v) \notin E$  then
  | | Add  $e'$  to  $E'$  with weight  $w'(e) = 1$ 
Compute the stationary distribution  $\pi'(v)$  of a random walk on  $G'$ 
 $G'' = (V, E'') \leftarrow G'$ 
for  $v \in V$  do
  | for  $e = (v, u) \in E''$  do
  | |  $w''(e) \leftarrow \frac{w'(e)}{\deg_{out}^+(v)} \pi'(v)$ 
Randomly sample an arborescence rooted at  $r$  from  $G''$ 

```

---

Thus, in RNC, the task of randomly sampling an arborescence in some digraph can be reduced to the task of randomly sampling an arborescence rooted at some vertex  $r$  from a strongly connected Eulerian digraph.

## 4 Cycle Decomposition

Because Theorem 5 only needs the first-visit edges to produce a random arborescence, our main idea is to create a hierarchical clustering of the graph that we call  $\mathcal{T}$ . Every element  $\mathcal{S} \in \mathcal{T}$  in the clustering is a subset of  $V$  which is strongly connected by edges of relatively high weight compared to the edges in  $\delta^-(\mathcal{S})$ . Intuitively, a random walk will spend much time traversing edges inside a cluster, before venturing out. Aside from first visit edges, any other traversals do not need to be simulated, thus much of the work can be avoided by only simulating the first few edges which enter or exit a cluster. Consider the pseudocode in Algorithm 2 for generating one such decomposition:

■ **Algorithm 2** Generating a hierarchical decomposition.

---

```

for edge  $e \in G$  in parallel do
  | Find a cycle  $C_e$  of edges such that  $e \in C_e$  and for every  $e' \in C_e$ ,  $w_{e'} \geq \frac{w_e}{m}$  using
  | Algorithm 3
Sort the cycles found by minimum edge weight in decreasing order  $C_1, C_2, \dots, C_m$ 
 $\mathcal{T} \leftarrow \{\}$ 
for  $i = 1, 2, \dots, m$  in parallel do
  | Find  $\mathcal{S}_{i,1}, \mathcal{S}_{i,2}, \dots, \mathcal{S}_{i,j}$ , the vertex sets of the connected components in
  |  $(V, C_1 \cup C_2 \cup \dots \cup C_i)$ 
  | Add all elements found to  $\mathcal{T}$ 
for  $\mathcal{S}_{i,j} \in \mathcal{T} \setminus \mathcal{S}_{m,1}$  in parallel do
  | Find the node of the form  $\mathcal{S}_{i+1,k}$  such that  $\mathcal{S}_{i,j} \subset \mathcal{S}_{i+1,k}$  and assign  $\mathcal{S}_{i+1,k}$  as the
  | parent of  $\mathcal{S}_{i,j}$ 
Contract duplicate nodes in  $\mathcal{T}$ 

```

---

## 83:10 Sampling Arborescences in Parallel

■ **Algorithm 3** FindCycle( $e = (e_0, e_1), G = (E, V)$ ).

---

```

Construct  $G' = (V, E')$ , where  $E' = \{e' \in E : w_{e'} \geq \frac{w_e}{m}\}$ 
for  $v_1, v_2 \in V$  in parallel do
  if  $(v_1, v_2) \in E'$  then
     $R(v_1, v_2, 1) \leftarrow \text{True}$ 
for  $v \in V$  in parallel do
   $R(v, v, 0) \leftarrow \text{True}$ 
   $R(v, v, 1) \leftarrow \text{True}$ 
for  $l \in \{2, 4, \dots, 2^{\lceil \ln n \rceil}\}$  do
  for  $v_1, v_2, v_3 \in V$  in parallel do
    if  $R(v_1, v_2, l/2)$  and  $R(v_2, v_3, l/2)$  then
       $R(v_1, v_3, l) \leftarrow \text{True}$ 
 $c_0 \leftarrow e_1, c_{2^{\lceil \ln n \rceil}} \leftarrow e_0$ 
for  $l \in \{2^{\lceil \ln n \rceil - 1}, 2^{\lceil \ln n \rceil - 2}, \dots, 1\}$  do
  for  $i \in [2^{\lceil \ln n \rceil}]$  such that  $l$  is the largest power of 2 which divides  $i$  in parallel do
    do
       $\text{Set } c_i \text{ to be some vertex such that } R(c_{i-l}, c_i, l) \text{ and } R(c_i, c_{i+l}, l)$ 
Prune the transcript of vertices  $\{c_0, c_1, \dots, c_{2^{\lceil \ln n \rceil}}\}$  to remove any duplicates and
return the resulting cycle.

```

---

We use Algorithm 3 to construct a cycle of comparably high-weight edges containing a given edge  $e$ . The collection of these cycles is a “loose decomposition” of the Eulerian graph into cycles. While an Eulerian graph can be decomposed into cycles in general, we do not know if this can be done in NC. Instead we find a collection of cycles whose sum and poly( $m$ ) times it sandwich the Eulerian graph.

For each edge  $e$ , Algorithm 3 successfully returns a cycle since,

► **Proposition 14.** *In directed weighted Eulerian graph  $G = (V, E)$  for any edge  $e \in E$ , there exists a cycle of edges  $C_e$  such that  $e \in C_e$  and for all  $e' \in C_e, w_{e'} \geq \frac{w_e}{m}$*

**Proof.** Note that any Eulerian digraph can be decomposed into at most  $m$  cycles where each cycle contains edges of uniform weight. Then, for any edge  $e$ , one of these cycles must have weight at least  $\frac{w_e}{m}$ . ◀

Given that a cycle is found containing every edge in  $G$ , it follows that Algorithm 2 successfully generates a hierarchical decomposition which satisfies:

► **Proposition 15.**  *$\mathcal{T}$  is a laminar family, i.e. has the following properties:*

1. Every node  $\mathcal{S}$  of  $\mathcal{T}$  is a subset of  $V$  and is the vertex set of a connected component of  $G$ .
2. Children of any node in  $\mathcal{T}$  are proper subsets of it.
3.  $|\mathcal{T}| \leq 2n$ .

Given this clustering  $\mathcal{T}$ , for any  $\mathcal{S} \in \mathcal{T}$ , we say an edge  $e \in E$  jumps between children of  $\mathcal{S}$  if  $\mathcal{S}$  is the lowest node in  $\mathcal{T}$  which contains both endpoints of  $e$ . We call  $e$  a *jumping edge* of  $\mathcal{S}$ .

In addition, we note that

► **Lemma 16.** *For any  $\mathcal{S} \in \mathcal{T}$ , let  $w_{\max}(\mathcal{S})$  be the maximum weight among the jumping edges of  $\mathcal{S}$ .  $G(\mathcal{S})$  is strongly connected by edges of weight  $\frac{w_{\max}(\mathcal{S})}{m}$ .*

**Proof.** We prove by contradiction. Assume that there is some edge  $e$  jumping between children of  $\mathcal{S}$  and two vertices  $u, v \in \mathcal{S}$  such that every path between  $u, v$  contains an edge of weight less than  $\frac{w_e}{m}$ . Let the cycles which connect  $\mathcal{S}$  in Algorithm 2 be  $C_1, \dots, C_k$ , with  $C_k$  being the cycle with an edge of minimum weight. Then, consider the path of edges connecting  $u$  to  $v$  only containing edges in these cycles. The weight of every edge in the path must be greater than  $w(C_k)$ , the weight of the minimum weight edge in  $C_k$ . Thus by our assumption we must have that  $\frac{w_e}{m} > w(C_k)$ .

But now consider  $C_e$ , the cycle found containing edge  $e$  with minimum edge weight  $w(C_e) \geq \frac{w_e}{m}$ . As  $w(C_e) > w(C_k)$ , this cycle must have been used to construct a child of  $\mathcal{S}$  in  $\mathcal{T}$ . This implies the  $e$  is contained in a child of  $\mathcal{S}$  and is therefore not a jumping edge of  $\mathcal{S}$ , a contradiction.  $\blacktriangleleft$

Lemma 16 intuitively states that the nodes in  $\mathcal{T}$  are tightly connected by edges of relatively high weight compared to incoming or outgoing edges. This will be critical in bounding the number of edges which need to be simulated.

## 5 Parallel Sampling

For clarity, a partially sampled random walk shall be represented as an alternating series of clusters and edges:  $\mathcal{S}_1, e_1, \mathcal{S}_2, e_2, \dots$ , where each  $\mathcal{S}_i \in \mathcal{T}$  and  $e_i \in \delta^-(\mathcal{S}_i) \cap \delta^+(\mathcal{S}_{i+1})$ . The sequence has a natural interpretation of a random walk in  $\mathcal{S}_i$  which then exits  $\mathcal{S}_i$  through  $e_i$  to continue in  $\mathcal{S}_{i+1}$  before traversing  $e_{i+1}$  and so on.

Once we have a hierarchical decomposition  $\mathcal{T}$  found in Section 4, a natural way to use such a decomposition is to simulate a random walk as follows:

1. Simulate just the edges jumping between high-level clusters of a random walk, producing a transcript of the form  $\mathcal{S}_1, e_1, \mathcal{S}_2, e_2, \dots$ .
2. For each  $i$ , recursively simulate a random walk on  $\mathcal{S}_i$  conditioned on exiting  $\mathcal{S}_i$  through edge  $e_i$ .

While this recursive procedure as described is not in RNC, using doubling tricks and memoization, we show how to perform this in RNC. One can view this as a generalization of the doubling trick used by Teng [31].

### 5.1 Simulating Jumping Edges

To simulate edges jumping across children of a node, we generalize the techniques of [31] to the following:

► **Theorem 17.** *Suppose that  $G = (V, E)$  is a directed graph, with weights given by  $w : E \rightarrow \mathbb{R}_{\geq 0}$ . Let  $\mathcal{S} \subseteq V$  be the disjoint union of  $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_k$ , where  $\mathcal{S}$  and each  $\mathcal{S}_i$  is strongly connected. The transcript of first  $L$  edges jumping across different  $\mathcal{S}_i$  in a random walk starting from a given vertex  $v_0 \in \mathcal{S}$  and conditioned on exiting  $\mathcal{S}$  through some edge  $e \in \delta^-(\mathcal{S})$  can be produced in the PRAM model using  $\text{poly}(|E|, L)$  processors and  $\text{polylog}(|E|, L)$  time.*

**Proof.** We denote this function  $\text{JumpingEdges}(\mathcal{S}, v_0, e, L)$ . See Algorithm 4 for pseudocode. Note that we can condition a walk on  $\mathcal{S}$  to exit via some edge  $e \in \delta^-(\mathcal{S})$  by considering a random walk on the graph induced by  $\mathcal{S}$  with an additional edge  $e$  and an absorbing state at the endpoint of  $e$ . For any vertex  $v \in \mathcal{S}$ , let  $\mathcal{S}_v$  denote the unique  $\mathcal{S}_i$  which contains  $v$ .

From Lemma 9, for any  $v \in \mathcal{S}$ , the probability that a random walk currently at  $v$  exits  $\mathcal{S}_v$  through each edge in  $\delta^-(\mathcal{S}_v)$  can be computed in NC. Thus, after using  $\text{poly}(|E|, L)$  processors and  $\text{polylog}(|E|, L)$  time to compute all such probabilities for every vertex  $v \in \mathcal{S}$ ,

## 83:12 Sampling Arborescences in Parallel

we can assume access to a deterministic function  $\text{Next}(v, X)$  which takes a vertex  $v \in \mathcal{S}$  and number  $X \in [0, 1]$  and outputs an edge  $e \in \delta^-(\mathcal{S}_v)$  such that if  $X$  is chosen uniformly from  $[0, 1]$ , then the probability that  $\text{Next}(v, X)$  returns some edge  $e \in \delta^-(\mathcal{S}_v)$  is the probability a random walk currently at  $v$  exits  $\mathcal{S}_v$  through edge  $e$ .

Given this function, if we generate independent uniformly random numbers  $X_1, \dots, X_L \in [0, 1]$ , the transcript of the jumping edges in random walk can be extracted by starting at  $v = v_0$  and repeatedly applying  $\text{Next}(v, X_i)$  to get the next vertex  $v$ , for  $i \in [L]$ . Of course, a naive implementation does not leverage parallelism, so we instead use a doubling trick.

We compute the values  $\text{End}(v, t, l)$  which would be the final edge jumping between children of  $\mathcal{S}$  when we run a random walk starting at  $v$ , and applying  $\text{Next}$  with random inputs  $X_t, X_{t+1}, \dots, X_{t+l-1}$ . For  $l = 1$ , these are 1-step random walks and we use  $\text{Next}$  to compute them. For simplicity, we allow the first argument of  $\text{End}$  to also be an edge, in which case the random walk starts at the ending vertex of that edge. Then we use the following identity, which allows us to compute  $\text{End}$  values for a particular  $l$  from  $\text{End}$  values for  $\frac{l}{2}$ .

$$\text{End}(v, t, l) = \text{End}(\text{End}(v, t, l/2), t, l/2)$$

This is because we can break an  $l$ -step random walk into two  $l/2$ -step random walks. So we can compute all  $\text{End}$  values when  $l$  ranges over powers of 2, in  $\log(L)$  steps.

Finally to compute the edge taken by the random walk at any time  $l = 1, \dots, L$ , we simply write down  $l$  as a sum of powers of 2, and repeatedly use the  $\text{End}$  function to compute the  $l$ -th edge of the random walk which jumps between children of  $\mathcal{S}$ . ◀

■ **Algorithm 4**  $\text{JumpingEdges}(\mathcal{S}, v_0, e_{\text{end}}, L)$ .

---

```

for  $t \in \{1, 2, \dots, L\}$  in parallel do
  Let  $X_t$  be a uniformly random number sampled from  $[0, 1]$ 
  for  $v \in \mathcal{S}$  in parallel do
     $\text{End}(v, t, 1) \leftarrow \text{Next}(v, X_t)$ 
for  $l \in \{2, 4, 8, \dots, 2^{\lfloor \ln L \rfloor}\}$  do
  for  $v \in \mathcal{S}, t \in \{0, 1, \dots, L - 1 - l\}$  in parallel do
     $\text{End}(v, t, l) \leftarrow \text{End}(\text{End}(v, t, \frac{l}{2}), t, \frac{l}{2})$ 
for  $l \in 1, 2, \dots, L$  in parallel do
  Decompose  $l$  as a sum of subset  $S \subset \{1, 2, \dots, 2^{\lfloor \ln L \rfloor - 1}\}$ 
   $e \leftarrow (*, v_0)$ 
   $t \leftarrow 0$ 
  for  $s \in S$  do
     $e \leftarrow \text{End}(e, t, s)$ 
     $t \leftarrow t + s$ 
   $e_l \leftarrow e$ 
   $v_l \leftarrow$  the ending vertex of  $e_l$ 
if  $e_k = e_{\text{end}}$  for some  $k$  then
   $\text{return } \mathcal{S}_{v_0}, e_1, \mathcal{S}_{v_1}, e_2, \dots, \mathcal{S}_{v_{k-1}}, e_k = e_{\text{end}}$ 
else
   $\text{return } \mathcal{S}_{v_0}, e_1, \mathcal{S}_{v_1}, e_2, \dots, \mathcal{S}_{v_{L-1}}, e_L, \mathcal{S}, e_{\text{end}}$ 

```

---

► **Remark 18.** For clarity of exposition, we allow sampling numbers from  $[0, 1]$ . However sampling a number with polynomially many bits is enough for our purposes. In our algorithms, we only need to compare our samples  $X_t$  with deterministic numbers derived from the input. To this end, we can simply sample the first  $N$  digits of the binary expansion of  $X_t$  for some large  $N$ . The probability that the comparison of  $X_t$  and a fixed number is not determined from the first  $N$  digits is  $2^{-N}$ . By taking  $N$  to be polynomially large, the probability of failure will be bounded by an exponentially small number. If one insists on avoiding even this small probability of failure, we can continue sampling digits of  $X_t$  whenever we run into a situation where the first  $N$  are not enough. Since the probability of running into failure is very small, the overall expected running time still remains polylogarithmic.

## 5.2 Extracting First-Visit Edges

For a cluster, once we know which edges jump across its children, we can recursively fill in the transcript of the walk. This is because we now have a similar subproblem for each child node, where we have a starting vertex, and a prespecified exit edge. However, a naive implementation of this is not in RNC.

Instead our strategy is to again use a doubling trick. Consider a node  $\mathcal{S}$  of  $\mathcal{T}$  and a vertex  $v \in \mathcal{S}$ . Using the algorithm `JumpingEdges`, we can extract the edges that jump across children of  $\mathcal{S}$ . We can then extend these to include edges that jump across children of children of  $\mathcal{S}$  by more applications of `JumpingEdges`. We will construct a function `AllEdges`( $\mathcal{S}, v, e, L, l$ ) that besides the arguments to `JumpingEdges` also takes an integer  $l \geq 1$ . The goal of this function is to extract from the transcript of a random walk started at  $v$  and conditioned on exiting  $\mathcal{S}$  through  $e$ , the first  $L$  edges which jump between descendants of depth at most  $l$  below  $\mathcal{S}$  in  $\mathcal{T}$ . Then, each intermediate node in the transcript is either an individual vertex, an  $l$ th descendant of  $\mathcal{S}$ , or a cluster for which  $L$  edges jumping between its children have already been found. We will show that for polynomially large  $L$ , this transcript will contain all first-visit edges with high probability.

A natural approach for computing `AllEdges` is as follows. We already know how to compute `AllEdges` for  $l = 1$ ; that is just a call to `JumpingEdges`. Once the value of `AllEdges` has been computed for all settings of parameters for a particular  $l$ , we can compute it for  $2l$  by the following procedure:

1. Let the transcript returned by `AllEdges`( $\mathcal{S}, v, e, L, l$ ) be  $\mathcal{S}_1, e_2, \mathcal{S}_2, e_2, \dots, \mathcal{S}_k, e_k$ .
2. For each intermediate node  $\mathcal{S}_i$  that is not a single vertex, if the transcript does not already contain  $L$  edges jumping between children of  $\mathcal{S}_i$ , let  $u$  be the endpoint of  $e_{i-1}$  and replace  $\mathcal{S}_i, e_i$  with `AllEdges`( $\mathcal{S}_i, u, e, L, l$ )
3. Trim the transcript to only contain at most the first  $L$  edges jumping between children of any node

There is one key problem with this approach. We cannot precompute `AllEdges`( $\cdot, \cdot, \cdot, \cdot, l$ ) and use the same precomputed answer to fill in the gaps for larger values of  $l$ . Each time we need to use `AllEdges` for a particular setting of its input parameters, we really need to use fresh randomness; otherwise the transcript we extract will not be from a true random walk.

We resolve this by using a caching trick. Instead of computing `AllEdges`( $\mathcal{S}, v, e, L, l$ ) for each setting of parameters once and reusing the same output for larger subproblems, we compute  $M$  possible answers for a large enough  $M$  and store all  $M$  answers. For larger subproblems, every time that we need to use the values of a smaller subproblem, we randomly pick one of the stored  $M$  answers. We will inevitably reuse some of the answers in this process; however when we restrict our attention to the unraveling of a particular answer for a subproblem there is a high probability of not having reused any answers. We will formalize this in Lemma 19.

## 83:14 Sampling Arborescences in Parallel

Pseudocode for computing `AllEdges` can be found in Algorithm 5. In the end we use `AllEdges(V, v,  $\emptyset$ , L, |V|)` to extract a list of edges, and with high probability all first-visit edges will be among this list

■ **Algorithm 5** Computing the random walk transcript extracts.

---

```

for  $\mathcal{S} \in \mathcal{T}, e \in \delta^-(\mathcal{S}) \cup \{\emptyset\}, v \in \mathcal{S}, i \in [M]$  in parallel do
  AllEdges( $\mathcal{S}, v, e, L, 1$ )[ $i$ ]  $\leftarrow$  JumpingEdges( $\mathcal{S}, v, e, L$ )
for  $l \in \{1, 2, \dots, 2^{\lceil \lg L \rceil - 1}\}$  do
  for  $\mathcal{S} \in \mathcal{T}, e \in \delta^-(\mathcal{S}) \cup \{\emptyset\}, v \in \mathcal{S}, i \in [M]$  in parallel do
    if  $l$  is greater than depth of the deepest child of  $\mathcal{S}$  then
      AllEdges( $\mathcal{S}, v, e, L, 2l$ )[ $i$ ]  $\leftarrow$  AllEdges( $\mathcal{S}, v, e, L, l$ )[ $i$ ]
    else
      Let  $\mathcal{S}_1, e_1, \mathcal{S}_2, e_2, \dots, \mathcal{S}_k$  be the output of AllEdges( $\mathcal{S}, v, e, L, l$ )[ $i$ ]
      for  $j \in [k]$  in parallel do
        if the transcript does not already contain  $L$  edges jumping between
          children of  $\mathcal{S}_j$  then
          Replace  $\mathcal{S}_j, e_j$  with a randomly chosen instance of
            AllEdges( $\mathcal{S}_j, v', e_j, L, l$ ) where  $v'$  is the endpoint of  $e_{j-1}$ 
          Trim the resulting transcript to only contain at most the first  $L$  edges
            jumping between children of any node. Save the resulting transcript as a
            solution to AllEdges( $\mathcal{S}, v, e, L, 2l$ )[ $i$ ]
      if the final transcript of AllEdges( $V, v, \emptyset, L, |V|$ ) contain multiple subsequences
        which depend on the same call to JumpingEdges then
        Replace all but the first subsequence with a freshly sampled transcript
  
```

---

► **Lemma 19.** *For  $M = \text{poly}(L)$ , when we unravel the recursion tree for the computation of the value `AllEdges`( $V, v, \emptyset, L, |V|$ ), no stored answer of `AllEdges` for any of the subproblems will be used more than once with high probability.*

**Proof.** Note that for any any  $\mathcal{S} \in \mathcal{T}, v \in \mathcal{S}, e \in \delta^-(\mathcal{S})$ , the recursion tree for an answer to `AllEdges`( $V, v, \emptyset, L, |V|$ ) will rely on at most  $L$  calls to `JumpingEdges`( $\mathcal{S}, v, e, L$ ) since each call to `JumpingEdges`( $\mathcal{S}, v, e, L$ ) corresponds to a subsequence of the final transcript which ends in  $e$ , and  $e$  can only occur  $L$  times in the transcript by the definition of `AllEdges`. As each sample of `JumpingEdges`( $\mathcal{S}, v, e, L$ ) is chosen uniformly at random from  $M$  samples, for polynomially large  $M$ , the probability of the same sample being chosen twice can be made small and of the form  $\frac{1}{\text{poly}(n)}$ . Lastly via a union bound over polynomially many possible combinations of arguments for `JumpingEdges`, we have that for polynomially large  $M$ , which high probability no stored answer of `JumpingEdges` or `AllEdges` will be used more than once. ◀

Putting everything together, we have shown that:

► **Theorem 20.** *Suppose that  $G = (V, E)$  is a directed graph, with weights given by  $w : E \rightarrow \mathbb{R}_{>0}$ . Let  $\mathcal{T}$  denote a hierarchical decomposition of  $G$ . The transcript of a random walk starting at  $v_0 \in V$  which contains first  $L$  edges jumping between children of  $\mathcal{S}$  for any  $\mathcal{S} \in \mathcal{T}$  can be produced in the PRAM model using  $\text{poly}(|E|, L)$  processors in  $\text{polylog}(|E|, L)$  expected time.*

## 6 Hierarchical Exploration Time

The last element needed to obtain an RNC algorithm for the random sampling of arborescences is to show that for the decomposition achieved above, for polynomially large  $L$ , simulating the first  $L$  edges which jump across each node  $S \in \mathcal{T}$  will contain all first-visit edges to each vertex with high probability. We do this by bounding the number of edge traversals between children of each node before that node is covered. We build on techniques developed by Boczkowski, Peres, and Sousi [5].

We start by bounding the number of times a given vertex is visited before cover time. For any  $v, s, t \in V$ , let  $H_v(s, t)$  denote the expected number of times a random walk starting at  $s$  visits  $v$  before reaching  $t$ ; in the case that  $s = t$ ,  $H_v(s, s)$  denotes the number of times a walk starting at  $s$  reaches  $v$  before returning to  $s$ .

It is easy to see that  $H_v$  satisfies a triangle inequality, namely:

► **Proposition 21.** *On any directed graph  $G = (V, E)$ , for any  $v, s, t, u, \in V$ :  $H_v(s, t) \leq H_v(s, u) + H_v(u, t)$ .*

Additionally, as the stationary distribution of vertices on an Eulerian graph is proportional to the degree of a vertex, it follows that:

► **Proposition 22.** *On any directed Eulerian graph  $G = (V, E)$ , for any  $v, s \in V$ :  $H_v(s, s) = \frac{1}{\pi(s)}\pi(v) = \frac{\deg(v)}{\deg(s)}$*

**Proof.** Consider the Schur complement  $G_S$  where  $S = \{v, s\}$ . Since  $G$  is Eulerian, so is  $G_S$  and the degrees of  $v, s$  are the same in  $G$  and  $G_S$ . Letting the weights of edges in  $G_S$  be  $w(v, s), w(s, v), w(v, v), w(s, s)$ , we must have  $w(s, v) = w(v, s)$  for  $G_S$  to be Eulerian. As random walks on  $G_S$  are distributed like random walks on  $G$  whose transcripts have been restricted to only contain vertices in  $S$ , the values of  $H_v(s, s), H_v(v, s)$  are the same for a walk on  $G$  as for one on  $G_S$ . Considering one-step transitions, we can construct the following system of equations:

$$\begin{aligned} H_v(s, s) &= \frac{w(s, v)}{\deg(s)} (H_v(v, s) + 1) + \frac{w(s, s)}{\deg(s)} 0, \\ H_v(v, s) &= \frac{w(v, v)}{\deg(v)} (H_v(v, s) + 1) + \frac{w(v, s)}{\deg(v)} 0. \end{aligned}$$

Solving yields  $H_v(v, s) = \frac{w(v, v)}{w(v, s)}$  and  $H_v(s, s) = \frac{w(s, v)}{\deg(s)} \left( \frac{w(v, v)}{w(v, s)} + 1 \right) = \frac{w(s, v)}{\deg(s)} \frac{\deg(v)}{w(v, s)} = \frac{\deg(v)}{\deg(s)}$ . ◀

For any two vertices which are connected by an edge, we have:

► **Lemma 23.** *On any directed Eulerian graph  $G = (V, E)$ , for any  $v, s, t \in V$  such that  $e = (s, t) \in E$ ,  $H_v(s, t) \leq \frac{\deg(v)}{w(s, t)}$*

**Proof.** Note that every time the walk is currently at  $s$ , there is a  $\frac{w(s, t)}{\deg(s)}$  chance of moving to  $t$ , thus, in expectation, the expected number of times the walk returns to  $s$  before reaching  $t$  is at most  $\frac{\deg(s)}{w(s, t)}$ . Between each return to  $s$ , the expected number of visits to  $v$  is  $\frac{\deg(v)}{\deg(s)}$  by Proposition 22. Thus, we conclude that:

$$H_v(s, t) \leq \frac{\deg(s)}{w(s, t)} \frac{\deg(v)}{\deg(s)} = \frac{\deg(v)}{w(s, t)}. \quad \blacktriangleleft$$

On graphs which are strongly connected by edges of some weight  $w_c$ , this observation leads to the conclusion that:

► **Lemma 24.** *On an Eulerian directed graph  $G = (V, E)$  strongly connected by edges of weight  $w_c$ , for any  $s, t \in V$ ,  $H_s(s, t) \leq \frac{\deg(s)n}{w_c}$*

**Proof.** By connectivity assumptions, there exists a sequence of at most  $n$  vertices  $s = v_0, v_1, \dots, v_k = t$  which form a path connecting  $s$  and  $t$  such that  $w(v_{i-1}, v_i) \geq w_c$  for all  $i \in [k]$ . Then, combining Proposition 21 with Lemma 23,

$$H_s(s, t) \leq \sum_{i=1}^k H_s(v_{i-1}, v_i) \leq \sum_{i=1}^k \frac{\deg(s)}{w(v_{i-1}, v_i)} \leq \frac{n \deg(s)}{w_c}. \quad \blacktriangleleft$$

This trivially bounds the expected number of returns to a vertex  $s$  before a graph  $G$  is covered by a random walk by  $\frac{n^2 \deg(s)}{w_c}$ .

► **Remark 25.** While this bound is sufficient for our purposes, using Matthew's trick, this bound can be further tightened to  $\frac{n \log n \deg(s)}{w_c}$  [26].

► **Lemma 26.** *For any node  $\mathcal{S} \in \mathcal{T}$  in the decomposition obtained by Algorithm 2, the expected number of edges traversed between children of  $\mathcal{S}$  before every vertex in  $\mathcal{S}$  has been visited is at most  $n^2 m^2$ .*

**Proof.** For each edge  $e = (u, v)$  which jumps between children of  $\mathcal{S}$ , the expected number of times edge  $e$  is traversed is  $\frac{w_e}{\deg(u)}$  times the expected number of times vertex  $u$  is reached by a random walk before all vertices in  $\mathcal{S}$  are reached. To bound the number of times  $u$  is reached before  $\mathcal{S}$  is covered, consider a random walk on the Schur complement  $G_{\mathcal{S}}$ . The expected number of times  $u$  is reached in a random walk on  $G$  before  $\mathcal{S}$  is covered is the same as the expected number of times  $u$  is reached by a random walk on  $G_{\mathcal{S}}$  before  $G_{\mathcal{S}}$  is covered.

By Proposition 11,  $G_{\mathcal{S}}$  will be Eulerian and the degrees of all vertices in  $\mathcal{S}$  will be the same in  $G$  and  $G_{\mathcal{S}}$ . By Lemma 16,  $\mathcal{S}$  is strongly connected by edges of weight  $\frac{w_{\max}(\mathcal{S})}{m}$ , and so  $G_{\mathcal{S}}$  is as well. Then, by Lemma 24, the expected number of times  $u$  is hit before  $G_{\mathcal{S}}$  is covered is at most by  $\frac{\deg(u)n^2 m}{w_{\max}}$ . Thus, the expected number of times edge  $e$  is traversed before  $\mathcal{S}$  is covered is at most

$$\frac{\deg(u)n^2 m}{w_{\max}} \frac{w_e}{\deg(u)} \leq n^2 m$$

since  $w_c \geq \frac{w_e}{m}$ . As there are at most  $m$  jumping edges of  $\mathcal{S}$ , in expectation, at most  $n^2 m^2$  edge traversals between children of  $\mathcal{S}$  will occur before every vertex in  $\mathcal{S}$  has been reached. ◀

Thus, for large enough  $L = \text{poly}(n, m)$ , by Markov's inequality, with high probability every node  $\mathcal{S} \in \mathcal{T}$  is covered by the time  $L$  edges have jumped across  $\mathcal{S}$ . This means the transcript returned by `AllEdges` will contain all first visit edges with high probability. As this transcript is polynomial in length, all first visit edges and the corresponding arborescence they form can be extracted and returned in NC.

## 7 Discussion and Open Problems

We showed how to sample arborescences, and as a special case spanning trees, from a given weighted graph using an RNC algorithm. While this is a step in resolving the disparity between parallel sampling and parallel counting algorithms, more investigation is needed. In particular, for the list of problems with determinant-based counting in Section 1, designing RNC sampling algorithms remains open.

One of these problems in particular, namely sampling Eulerian tours from an Eulerian digraph, is intimately connected to sampling arborescences, due to the BEST theorem [1, 34]. However the reduction, while polynomial-time implementable, is not known to be in NC.



► **Question 27.** Is there an algorithm for sampling Eulerian tours uniformly at random in Eulerian digraphs?

The important tool our result relied on was the Aldous-Broder algorithm. So it is natural to ask whether there are generalizations of the Aldous-Broder result to settings beyond spanning trees and arborescences. In particular the bases of a *regular matroid* are a proper generalization of spanning trees in a graph, and they have a decomposition in terms of graphic, co-graphic, and some special constant-sized matroids [29].

► **Question 28.** Can we sample from a regular matroid, or equivalently from a volume-based distribution defined by totally unimodular vectors in RNC?

Yet another direction for generalization are higher-dimensional equivalents of spanning trees and arborescences. For example, Gorodezky and Pak [15] provided a generalization of the algorithm of Wilson [35] for sampling arborescences from graphs to hypergraphs. While Wilson’s algorithm is different from that of Aldous-Broder, it is closely related. Can these generalizations be efficiently parallelized?

---

## References

- 1 T Aardenne-Ehrenfest and NG Bruijn. Circuits and trees in oriented linear graphs. *Classic Papers in Combinatorics*, pages 149–163, 1987.
- 2 David J Aldous. The random walk construction of uniform spanning trees and uniform labelled trees. *SIAM Journal on Discrete Mathematics*, 3.4:450–465, 1990.
- 3 Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials iv: Exchange properties, tight mixing times, and faster sampling of spanning trees. *arXiv preprint arXiv:2004.07220*, 2020.
- 4 Nikhil Balaji and Samir Datta. Tree-width and logspace: Determinants and counting euler tours. *arXiv preprint*, 2013. [arXiv:1312.7468](https://arxiv.org/abs/1312.7468).
- 5 Lucas Boczkowski, Yuval Peres, and Perla Sousi. Sensitivity of mixing times in eulerian digraphs. *SIAM Journal on Discrete Mathematics*, 32(1):624–655, 2018.
- 6 Andrei Broder. Generating random spanning trees. In *Foundations of Computer Science, 1989., 30th Annual Symposium on*, pages 442–447. IEEE, 1989.
- 7 Charles J Colbourn, Wendy J Myrvold, and Eugene Neufeld. Two algorithms for unranking arborescences. *Journal of Algorithms*, 20(2):268–281, 1996.
- 8 Patrick John Creed. *Counting and sampling problems on Eulerian graphs*. PhD thesis, The University of Edinburgh, 2010.
- 9 Laszlo Csanky. Fast parallel matrix inversion algorithms. In *Foundations of Computer Science, 1975., 16th Annual Symposium on*, pages 11–12. IEEE, 1975.
- 10 Amit Deshpande and Luis Rademacher. Efficient volume sampling for row/column subset selection. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 329–338. IEEE, 2010.
- 11 David Durfee, Rasmus Kyng, John Peebles, Anup B Rao, and Sushant Sachdeva. Sampling random spanning trees faster than matrix multiplication. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 730–742, 2017.
- 12 David Durfee, John Peebles, Richard Peng, and Anup B Rao. Determinant-preserving sparsification of sddm matrices with applications to counting and sampling spanning trees. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 926–937. IEEE, 2017.
- 13 Martin Dyer, Alan Frieze, and Ravi Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies. *Journal of the ACM (JACM)*, 38(1):1–17, 1991.
- 14 Ira M Gessel and Xavier Viennot. Determinants, paths, and plane partitions. *preprint*, 132(197.15), 1989.

- 15 Igor Gorodezky and Igor Pak. Generalized loop-erased random walks and approximate reachability. *Random Structures & Algorithms*, 44(2):201–223, 2014.
- 16 Mark Jerrum and Alistair Sinclair. The markov chain monte carlo method: an approach to approximate counting and integration. *Approximation algorithms for NP-hard problems*, pages 482–520, 1996.
- 17 Mark Jerrum, Alistair Sinclair, and Eric Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. *Journal of the ACM (JACM)*, 51(4):671–697, 2004.
- 18 Mark R Jerrum, Leslie G Valiant, and Vijay V Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical computer science*, 43:169–188, 1986.
- 19 Minghui Jiang, James Anderson, Joel Gillespie, and Martin Mayne. ushuffle: a useful tool for shuffling biological sequences while preserving the k-let counts. *BMC bioinformatics*, 9(1):192, 2008.
- 20 Pieter W Kasteleyn. Dimer statistics and phase transitions. *Journal of Mathematical Physics*, 4(2):287–293, 1963.
- 21 Jonathan A Kelner and Aleksander Madry. Faster generation of random spanning trees. In *Foundations of Computer Science, 2009. FOCIS'09. 50th Annual IEEE Symposium on*, pages 13–21. IEEE, 2009.
- 22 Gustav Kirchhoff. Ueber die auflösung der gleichungen, auf welche man bei der untersuchung der linearen vertheilung galvanischer ströme geführt wird. *Annalen der Physik*, 148(12):497–508, 1847.
- 23 David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.
- 24 Bernt Lindström. On the vector representations of induced matroids. *Bulletin of the London Mathematical Society*, 5(1):85–90, 1973.
- 25 Aleksander Mađry, Damian Straszak, and Jakub Tarnawski. Fast generation of random spanning trees and the effective resistance metric. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 2019–2036. Society for Industrial and Applied Mathematics, 2015.
- 26 Peter Matthews. Covering problems for brownian motion on spheres. *The Annals of Probability*, 16(1):189–199, 1988. doi:10.1214/aop/1176991894.
- 27 Romain Rivière, Dominique Barth, Johanne Cohen, and Alain Denise. Shuffling biological sequences with motif constraints. *Journal of Discrete Algorithms*, 6(2):192–204, 2008.
- 28 Aaron Schild. An almost-linear time algorithm for uniform random spanning tree generation. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 214–227. ACM, 2018.
- 29 Paul D Seymour. Decomposition of regular matroids. *Journal of combinatorial theory, Series B*, 28(3):305–359, 1980.
- 30 Shang-Hua Teng. Independent sets versus perfect matchings. *Theoretical Computer Science*, 145(1-2):381–390, 1995.
- 31 Shang-Hua Teng. Independent sets versus perfect matchings. *Theoretical Computer Science*, 145.1-2:381–390, 1995.
- 32 Prasad Tetali and Santosh Vempala. Random sampling of euler tours. *Algorithmica*, 30(3):376–385, 2001.
- 33 W. T. Tutte. The dissection of equilateral triangles into equilateral triangles. *Mathematical Proceedings of the Cambridge Philosophical Society*, 44(4):463–482, 1948. doi:10.1017/S030500410002449X.
- 34 WT Tutte and Cedric AB Smith. On unicursal paths in a network of degree 4. *The American Mathematical Monthly*, 48(4):233–237, 1941.
- 35 David Bruce Wilson. Generating random spanning trees more quickly than the cover time. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 296–303, 1996.