

Sink-Free Orientations: A Local Sampler with Applications

Konrad Anand 

University of Edinburgh, UK

Graham Freifeld 

University of Edinburgh, UK

Heng Guo 

University of Edinburgh, UK

Chunyang Wang 

Nanjing University, China

Jiaheng Wang 

University of Regensburg, Germany

Abstract

For sink-free orientations in graphs of minimum degree at least 3, we show that there is a deterministic approximate counting algorithm that runs in time $O((n^{33}/\varepsilon^{32}) \log(n/\varepsilon))$, a near-linear time sampling algorithm, and a randomised approximate counting algorithm that runs in time $O((n/\varepsilon)^2 \log(n/\varepsilon))$, where n denotes the number of vertices of the input graph and $0 < \varepsilon < 1$ is the desired accuracy. All three algorithms are based on a local implementation of the sink popping method (Cohn, Pemantle, and Propp, 2002) under the partial rejection sampling framework (Guo, Jerrum, and Liu, 2019).

2012 ACM Subject Classification Theory of computation → Random walks and Markov chains

Keywords and phrases Sink-free orientations, local sampling, deterministic counting

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2025.60

Category RANDOM

Related Version *arXiv version*: <https://arxiv.org/abs/2502.05877>

Funding This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 947778). Jiaheng Wang also acknowledges support from the ERC (grant agreement No. 101077083, CountHom).

Acknowledgements We thank Guus Regts for helpful comments on an earlier version of this paper.

1 Introduction

The significance of counting has been recognised in the theory of computing since the pioneering work of Valiant [44, 43]. In the late 80s, a number of landmark approximate counting algorithms [30, 18, 31] were discovered. A common ingredient of these algorithms is the computational equivalence between approximate counting and sampling for self-reducible problems [32]. The reduction from counting to sampling decomposes the task into a sequence of marginal probability estimations, each of which is tractable for sampling techniques such as Markov chains. However, while only the marginal probability of one variable is in question, simulating Markov chains requires keeping track of the whole state of the instance, and thus is obviously wasteful. It is more desirable to draw samples while accessing only some local structure of the target variable. We call such algorithms local samplers.

The first such local sampler was found by Anand and Jerrum [2], who showed how to efficiently generate perfect local samples for spin systems even when the underlying graph is infinite. Using local information is essential here as it is not possible to perfectly simulate



© Konrad Anand, Graham Freifeld, Heng Guo, Chunyang Wang, and Jiaheng Wang; licensed under Creative Commons License CC-BY 4.0
Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2025).

Editors: Alina Ene and Eshan Chattopadhyay; Article No. 60; pp. 60:1–60:19



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

the whole state. Subsequently, Feng, Guo, Wang, Wang, and Yin [20] found an alternative local sampler, namely the so-called coupling towards the past method, which yields local implementations of rapid mixing Markov chains. It is also observed that sufficiently efficient local samplers lead to immediate derandomisation via brute-force enumeration. Moreover, local samplers are crucial to obtain sub-quadratic time approximate counting algorithms for spin systems [1]. Thus, local samplers are highly desirable algorithms as they can lead to fast sampling, fast approximate counting, and deterministic approximate counting algorithms.

Guo, Jerrum, and Liu [25] introduced partial rejection sampling (PRS) as yet another efficient sampling technique. This method generalises the so-called cycle-popping algorithm for sampling spanning trees [46] and the sink-popping algorithm for sampling sink-free orientations [16]. It also has close connections with the Lovász local lemma [19]. For extremal instances (in the sense of [33]), PRS is just the celebrated Moser-Tardos algorithm for the constructive local lemma [39]. The most notable application of PRS is the first fully polynomial-time randomised approximation scheme (FPRAS) for all-terminal network reliability [23]. On the other hand, it is still open if all-terminal reliability and counting sink-free orientations admit deterministic fully polynomial-time approximation schemes (FPTASes). Thus, in view of the aforementioned derandomisation technique [20], a local implementation of PRS is a promising way to resolve these open problems.

In this paper, we make some positive progress for sink-free orientations (SFOs). Given an undirected graph $G = (V, E)$, a sink-free orientation of G is an orientation of its edges such that each vertex has at least one outgoing edge. SFOs were first studied by Bubley and Dyer [10] as a restricted case of SAT.¹ They showed that exact counting of SFOs is $\#\mathbf{P}$ -complete, and thus is unlikely to have a polynomial-time algorithm. For approximate counting and sampling, in [11], they showed that a natural Markov chain has an $O(m^3)$ mixing time, where m is the number of edges. Later, Cohn, Pemantle, and Propp [16] introduced an exact sampler, namely the aforementioned sink-popping algorithm that runs in $O(nm)$ time in expectation, where n is the number of vertices. Using the partial rejection sampling framework, Guo and He [22] improved the running time of sink-popping to $O(n^2)$, and constructed instances where this running time is tight. It is open whether a faster sampling algorithm or an FPTAS exists.

Our main result is a local sampler based on partial rejection sampling for SFOs. Using this local sampler, for graphs of minimum degree 3, we obtain a deterministic approximate counting algorithm that runs in time $O((n^{33}/\varepsilon^{32}) \log(n/\varepsilon))$, a near-linear time sampling algorithm, and a randomised approximate counting algorithm that runs in time $O((n/\varepsilon)^2 \log(n/\varepsilon))$, where ε is the given accuracy. All three algorithms appear difficult to obtain using previous techniques. We will describe the results in more detail in the next section.

1.1 Our contribution and technique overview

Our local sampler works for a slight generalisation of SFOs, which are intermediate problems required by the standard counting to sampling reduction [32]. In these problems, a subset S of vertices is specified, which are required to be sink-free, and the task is to estimate the probability of a vertex v not in S not being a sink. We denote by μ_S the uniform distribution of orientations where no vertex in S is a sink.

¹ We remark that a variant of SFOs is also introduced in the context of distributed computing under the name of sinkless orientations, where they are used to give a lower bound for the distributed Lovász local lemma [8]. The main difference is that for sinkless orientations, usually only vertices of sufficiently high degrees are required to be non-sinks.

Before describing our technique, let us first review the sink-popping algorithm, (which is a special case of partial rejection sampling and the same as the Moser-Tardos algorithm [39] as the instance is extremal). We orient each edge uniformly at random. As long as there is a sink in S , we select one such vertex, arbitrarily, and rerandomise all edges incident to it, until there is no sink belonging to S .

Our key observation is that it is unnecessary to simulate all edges to decide if $v \notin S$ is a sink. In particular, if, at any point of the execution of the algorithm, v is a sink, then no adjacent edges will ever be resampled and v stays a sink till the algorithm finishes. On the other hand, if at any point v , using only the already oriented edges, belongs to a cycle, a path leading to a cycle, or a nonempty path leading to some vertex not in S , then the orientations of all edges involved will not be resampled, and v stays a non-sink until the algorithm terminates. Thus, this observation gives us an early termination criterion for determining whether v is a sink or not. Moreover, since in the sink-popping algorithm, the order of sinks popped can be arbitrary, we can reveal the random orientation of edges strategically, and pop sinks if necessary. To be more precise, we first reveal the edges adjacent to v one by one. Once there is an outgoing edge (v, u) , we then move to u and repeat this process. If any sink is revealed, we erase the orientations of all its adjacent edges and backtrack. Eventually, one of the two early termination rules above will kick in, and this gives us a local sample.

Ideally, we want our local sampler to run in $O(\log n)$ time, where n is the number of vertices. Unfortunately, the one described above does not necessarily terminate this fast. To see this, consider a sequence of degree 2 vertices, where at each step there is equal probability to move forward or backtrack. Resolving such a path of length ℓ would require $\Theta(\ell^2)$ time. On the other hand, when the minimum degree of the input graph is at least 3, the length of the path followed by the sampler forms a submartingale. The vertex v can be a sink only if this path has length 0. Thus, once the length of the path is at least $C \log n$ for some constant C , the probability of v being a sink is very small. This allows us to truncate the local sampler with only a small error. Recall that μ_S is the uniform distribution of orientations under which no vertex in S is a sink. Our main theorem is as follows.

► **Theorem 1** (local sampler). *There exists an algorithm that, given a graph $G = (V, E)$ with minimum degree at least 3, $S \subseteq V$, and $v \notin S$, samples a Boolean random variable x in time $O(\log \frac{1}{\varepsilon})$ such that $|\Pr[x = 1] - \mu_S(v \text{ is not a sink})| \leq \varepsilon$.*

Moreover, there exists an algorithm that, in the same setting, given any edge $e \in E$, samples a random orientation of e that is ε -close in total variation distance from the marginal distribution of e under μ_S .

We note that for the algorithm in Theorem 1, we assume that the input graph can be accessed by an adjacent neighbour oracle. That is, given a vertex v and a number k , it returns the k -th neighbour of v . Also, all our algorithms work for not necessarily simple graphs.

The FPTAS for #SFO follows from the derandomisation method of [20] to the truncated local sampler. By ε -approximation, we mean an estimate \tilde{Z} such that $1 - \varepsilon \leq \frac{\tilde{Z}}{Z} \leq 1 + \varepsilon$, where Z is the target quantity.

► **Corollary 2** (deterministic approximate counting). *For graphs with minimum degree at least 3, there exists a deterministic algorithm that, given $0 < \varepsilon < 1$, outputs an ε -approximation to the number of sink-free orientations with running time $O((n^{33}/\varepsilon^{32}) \log(n/\varepsilon))$, where n is the number of vertices.*

Although high, the constant exponent in the running time of Corollary 2 is actually the most interesting feature of our algorithm. This exponent is derived from the hidden constant in Theorem 1. In contrast, the running time of most known FPTASes [45, 4, 6, 28, 5, 40,

[37, 20, 12] has an exponent that depends on some parameter (such as the maximum degree) of the input graph. There are exceptions, for example, [34, 27], but the exponents of their running times still depend on the parameters of the problem (not of the instance).

Another corollary of Theorem 1 is the following fast approximate sampler.

► **Corollary 3** (fast sampling). *For graphs with minimum degree at least 3, there exists a sampling algorithm that, given $0 < \varepsilon < 1$, outputs a random orientation σ such that σ is ε -close to a uniform random sink-free orientation in total variation distance, with running time $O\left(m \log\left(\frac{m}{\varepsilon}\right)\right)$, where m is the number of edges.*

We note that when the minimum degree is at least 4, the same running time can be achieved by the Moser-Tardos algorithm [38], as the symmetric local lemma condition holds in that case. However, when the minimum degree is 3, the symmetric local lemma condition no longer holds and our analysis relies on Shearer’s condition [42]. Under Shearer’s condition, we are not aware of a similar running time bound of the MT algorithm.

Our sampler runs in $\tilde{O}(m)$ time² instead of the $O(n^2)$ time that sink-popping requires, at the cost of generating an approximate sample instead of a perfect sample. This improves over sink-popping when $m = o(n^2/\log n)$ and leads to a faster FPRAS using the counting-to-sampling reduction [32]. In fact, the running time of the FPRAS can be improved further by directly invoking the truncated local sampler in the reduction.

► **Corollary 4** (fast approximate counting). *For graphs with minimum degree at least 3, there exists a (randomised) algorithm that, given $0 < \varepsilon < 1$, outputs a quantity that is an ε -approximation with probability at least $3/4$ to the number of sink-free orientations. The running time is $O((n/\varepsilon)^2 \log(n/\varepsilon))$, where n is the number of vertices.*

The success probability $3/4$ in Corollary 4 is standard in the definition of FPRAS, and can be easily amplified by taking the median of repeated trials and applying the Chernoff bound.

Note that directly combining Corollary 3 with the counting-to-sampling reduction results in an $\tilde{O}(nm/\varepsilon^2)$ running time. Corollary 4 is faster when $m = \omega(n)$. Previously, the best running time for approximate counting is $\tilde{O}(n^3/\varepsilon^2)$, via combining the $O(n^2)$ time sink-popping algorithm [22] with simulated annealing (see, for example, [22, Lemma 12]). Corollary 4 improves over this by roughly a factor of n . In very dense graphs (when $m = \Omega(n^2)$), Corollary 4 achieves near-linear time, which appears to be rare for approximate counting.

1.2 Related work

Our local sampler in Theorem 1 also falls under the framework of local access of huge random objects, a question first proposed by Goldreich, Goldwasser, and Nussboim [21]. This is a sampling version of local computation algorithms (LCAs) [41]. As such, the LCA lower bound for sinkless orientations by Brandt, Grunau, and Rozhon [9] implies that the running time of Theorem 1 is nearly optimal. Moreover, an early termination implementation of the general Moser-Tardos algorithm was developed by Davies-Peck [17] for LCAs and distributed algorithms. However, the algorithm in [17] requires a “polynomially-weakened” symmetric local lemma condition, whereas in the setting of Theorem 1 only Shearer’s condition for the local lemma holds. In addition, the algorithm in [17] appears to be too slow for our derandomisation purposes.

² The \tilde{O} notation hides logarithmic factors.

There are a plethora of fast sampling and deterministic approximate counting techniques by now. However, it appears difficult to achieve our results without the new local sampler. For example, the coupling of Bubley and Dyer [11] does not seem to improve with the minimum degree requirement. On a similar note, the recent deterministic counting technique of [12] requires a distance-decreasing Markov chain coupling, whereas the Bubley-Dyer coupling is distance non-increasing. In any case, even if the technique of [12] applied, it would not imply a running time with a constant exponent. Other fast sampling and FPTAS techniques, such as spectral independence [3, 15, 14], correlation decay [45, 35], and zero-freeness of polynomials [5, 40, 26], all seem difficult to apply. The main obstacle is that these techniques typically make use of properties that hold under arbitrary conditionings. However, for SFO, even if we start with a graph of minimum degree 3, conditioning the edges can result in a graph that is effectively a cycle, in which case no nice property holds. Our techniques, in contrast, require no hereditary properties and thus can benefit from the minimum degree requirement.

One much less obvious alternative approach to FPTAS is via the connection of the local lemma. In particular, because SFOs form extremal instances, their number can be computed via the independence polynomial evaluated at negative weights on the dependency graph. (We also see this fact in Section 4.1.) Normally this approach would not be efficient, because the dependency graph is usually exponentially large (for example for all-terminal reliability), but in the case of SFOs, the dependency graph is just the input graph itself. There are more than one FPTASes [40, 28] for the independence polynomial at negative weights. However, neither appears able to recover Corollary 2. With the minimum degree ≥ 3 assumption, the probability vector for SFOs is within the so-called Shearer’s region, where both algorithms apply.³ The downside is that the running time of both algorithms has the form $(n/\varepsilon)^{O(\log d)}$,⁴ where d is the maximum degree of the graph. Thus, in the setting of Corollary 2, these algorithms run in quasi-polynomial time instead. A more detailed discussion is given in Section 4.2.

Very recently, after our paper was posted on arXiv, Bencs and Regts [7] showed that in fact, the connection above can lead to an FPTAS for SFOs via Barvinok’s interpolation method. The main novel idea is that one need to instantiate the independence polynomial in a different way. As in Corollary 2, their algorithm also requires minimum degree ≥ 3 , but their running time is faster, being $O(n(m/\varepsilon)^{6.47})$, where n is the number of vertices and m is the number of edges of the graph.

1.3 Organisation

The rest of the paper is organised as follows. In Section 2, we introduce our local sampler. It is then analysed in Section 3. The main theorems are shown in Section 4. We conclude with a few open problems in Section 5.

³ In [40], only a uniform bound is stated, but one can introduce a scaling variable t and make a new polynomial in t , so that their algorithm works in the Shearer’s region.

⁴ To be more precise, the hidden constants in the exponents decrease in the multiplicative “slack” of how close the evaluated point is to the boundary of Shearer’s region. For SFOs, when constant degree vertices are present, the slack is a constant, and so are the hidden constants in the exponents.

2 A local sampler for sink-free orientations

Fix $G = (V, E)$ as an undirected graph. An orientation σ of G is an assignment of a direction to each edge, turning the initial graph into a directed graph. For any $S \subseteq V$, let Ω_S be the set of S -sink-free orientations of G , i.e., the set of orientations such that each vertex $v \in S$ is not a sink. Thus, Ω_V is the set of all (normal) sink-free orientations of G . When $|\Omega_S| \neq 0$, we use μ_S to denote the uniform distribution over Ω_S . For two adjacent vertices $u, v \in V$, we use $\{u, v\}$ to denote the undirected edge and (u, v) to denote the directed edge, from u to v .

We apply the following standard counting-to-sampling reduction [32]. Let $V = \{v_1, v_2, \dots, v_n\}$ be arbitrarily ordered and, for each $0 \leq i \leq n$, define $V_i = \{v_1, v_2, \dots, v_i\}$. Then, $|\Omega_V|$ can be decomposed into a telescopic product of marginal probabilities:

$$|\Omega_V| = |\Omega_{V_0}| \cdot \prod_{i=1}^n \frac{|\Omega_{V_i}|}{|\Omega_{V_{i-1}}|} = 2^{|E|} \cdot \prod_{i=1}^n \mu_{V_{i-1}}(v_i \text{ is not a sink}). \quad (1)$$

Thus, our goal becomes to estimate $\mu_S(v \text{ is not a sink})$ for any $S \subseteq V$ and $v \notin S$.

We view S -sink-free orientations under the variable framework of the Lovász local lemma. Here, each edge corresponds to a variable that indicates its direction, and each vertex in S represents a bad event of being a sink. An instance is called *extremal* if any two bad events are independent (namely, they share no common variable) or disjoint. It is easy to see that all instances to the S -sink-free orientation problem are extremal: if a vertex is a sink then none of its neighbors can be a sink. For extremal instances like this, the celebrated Moser-Tardos algorithm [39] is guaranteed to output an assignment avoiding all bad events uniformly at random [25]. This is summarised in Algorithm 1. Note that when $S = V$, Algorithm 1 is the sink-popping algorithm by Cohn, Pemantle, and Propp [16].

■ **Algorithm 1** Partial rejection sampling for generating an S -sink-free orientation.

Input : an undirected graph $G = (V, E)$ and a subset of vertices $S \subseteq V$
Output : an orientation σ of G


- 1 orient each edge $e \in E$ uniformly at random and independently to obtain an orientation σ ;
- 2 **while** $\exists v \in S$ s.t. v is a sink in σ **do**
- 3 choose such a v arbitrarily;
- 4 resample the orientation of all edges incident to v in σ uniformly at random;
- 5 **return** σ ;

The following lemma is a direct corollary from [25, Theorem 8] and SFOs being extremal.

► **Lemma 5.** *If $|\Omega_S| \neq 0$, Algorithm 1 terminates almost surely and returns an orientation distributed exactly as μ_S .*

We remark that the only possible case for $\Omega_S = \emptyset$ is when S forms a tree and is not connected to any vertex not in S .

Algorithm 1 requires one to generate a global sample when estimating $\mu_S(v \text{ is not a sink})$ for some $v \notin S$, which is wasteful. The following observation is crucial to turning it into a local sampler.



■ **Figure 1** Illustration of Lemma 6. Shaded vertices are in the set S . Once these patterns are formed, thick red edges would never be resampled in Algorithm 1.

► **Lemma 6** (criteria for early termination). *Suppose $|\Omega_S| \neq 0$. For any $v \notin S$, v is a sink upon the termination of Algorithm 1 if and only if*

(a) *v becomes a sink at some iteration.*

Conversely, v is not a sink upon the termination of Algorithm 1 if and only if one of the following holds:

(b1) *a directed cycle C containing v , or a directed path P containing v which ends in a directed cycle C is formed in some iteration, or*

(b2) *a nonempty directed path P from v to some $u \notin S$ is formed in some iteration.*

Proof. First, consider (a). If v becomes a sink at any point, then for every $w \in S$ which is a neighbour of v , the edge (w, v) is oriented towards v . Since $v \notin S$, the edge (w, v) will not be resampled via v , and could only be resampled by w becoming a sink. Since w cannot become a sink without resampling (w, v) , v will remain a sink. The other implication is obvious.

Now for (b1) and (b2), we first consider the forward implications. For a cycle C , every vertex $u \in C$ has an edge pointing outwards towards some $w \in C$ which also has an edge pointing outwards. None of these edges can be resampled without another edge $e \in C$ being resampled first, so no vertex in the cycle will ever become a sink again.

Consider a path P that ends outside of S or in a cycle. Inductively, we see that no edge on this path will be resampled without the edge after it being resampled. The edge connected to the cycle or S^c will not be resampled, since that vertex may never be a sink, so no vertex in P can become a sink.

For the reverse implication, suppose v is eventually not a sink. In that case, there must be some edge (v, w) pointing towards a neighbour w . If w is a sink, then $w \notin S$ and we are in case (b2). Otherwise, w is not a sink, and there is an adjacent edge pointing away from w , and we repeat this process. As the set of vertices is finite, if vertices considered in this process are all in S , then there must be repeated vertices eventually, in which case we are in (b1). ◀

An illustration of Lemma 6 is given in Figure 1. Based on Lemma 6, we design a local sampling algorithm for determining whether some vertex $v \notin S$ is a sink or not, given in Algorithm 2. We assume that the undirected graph $G = (V, E)$ is stored as an adjacency list where the neighbors of each vertex are arbitrarily ordered. Algorithm 2 takes as input some $S \subseteq V$ and $v \notin S$, returns an indicator variable $x \in \{0, 1\}$ such that $\Pr[x = 1] = \mu_S(v \text{ is not a sink})$. We treat the path P as a subgraph, and $V(P)$ denotes the vertex set of P . When we remove the last vertex from P , we remove it and the adjacent edge as well. Informally, Algorithm 2 starts from the vertex v , and reveals adjacent edges one by one. If there is any edge pointing outward, say (v, u) , we move to u and reveal edges adjacent to u . If a sink $w \in S$ is formed, then we mark all adjacent edges of w as unvisited

and backtrack. This induces a directed path starting from v . We repeat this process until any of the early termination criteria of Lemma 6 is satisfied, in which case we halt and output accordingly.

■ **Algorithm 2** $\text{Sample}(G, S, v)$.

Input : an undirected graph $G = (V, E)$, a subset of vertices $S \subseteq V$, and a vertex $v \notin S$;

Output : a random value $x \in \{0, 1\}$;

- 1 Let P be a (directed) path initialised as $P = (v)$;
- 2 Initialise a mapping $M : E \rightarrow \{\text{visited}, \text{unvisited}\}$ so that $\forall e \in E, M(e) = \text{unvisited}$;
- 3 **while** $|V(P)| \geq 1$ **do**
- 4 Let u be the last vertex of P ;
- 5 **if** $|V(P)| \geq 2$ *and* $u \notin S$ **then return** 1;
- 6 **if** *all edges incident to u are marked visited* **then**
- 7 mark all edges incident to u as unvisited;
- 8 remove u from P ;
- 9 **else**
- 10 let $e = \{u, w\}$ be the first unvisited edge incident to u ;
- 11 mark e as visited;
- 12 **with probability** $1/2$ **do**
- 13 **if** $w \in V(P)$, *or there is a visited edge (w, w') for some $w' \in V(P)$* **then**
- 14 **return** 1;
- 14 append w to the end of P ;
- 15 **return** 0;

3 Analysis of the local sampler

In this section, we analyse the correctness and efficiency of Algorithm 2.

► **Lemma 7** (correctness of Algorithm 2). *Let $G = (V, E)$ be a graph. For any $S \subseteq V$ such that $|\Omega_S| \neq 0$ and any $v \notin S$, $\text{Sample}(G, S, v)$ terminates almost surely and upon termination, returns an $x \in \{0, 1\}$ such that*

$$\Pr[x = 1] = \mu_S(v \text{ is not a sink}).$$

Proof. We claim that there exists a coupling between the execution of Algorithm 1 (with input G, S and output σ), and $\text{Sample}(G, S, v)$ (with output x), such that

$$v \text{ is not a sink under } \sigma \iff x = 1. \quad (2)$$

The claim implies the lemma because of Lemma 5.

To prove the claim, we first construct our coupling. We use the resampling table idea of Moser and Tardos [39]. For each edge, we associate it with an infinite sequence of independent random variables, each of which is a uniform orientation. This forms the “resampling table”. Our coupling uses the same resampling table for both Algorithm 1 and Algorithm 2. As showed in [39], the execution of Algorithm 1 can be viewed as first constructing this (imaginary) table, and whenever the orientation of an edge is randomised, we just reveal and use the next random variable in the sequence. For Algorithm 2, we reveal

the orientation of an edge when its status changes from *unvisited* to *visited* in Line 11. We execute Line 14 if the random orientation from the resampling table is (u, w) , and otherwise do nothing. Namely, in the latter case, the revealed orientation is (w, u) , and we just move forward the “frontier” of that edge by one step in the resampling table. We claim that (2) holds with this coupling.

Essentially, the claim holds since for extremal instances, given a resampling table, the order of the bad events resampled in Algorithm 1 does not affect the output. This fact is shown in [24, Section 4]. (See also [16, Lemma 2.2] for the case of $S = V$.) We can “complete” Algorithm 2 after it finishes. Namely, once Algorithm 2 terminates, we randomise all edges that are not yet oriented, and resample edges adjacent to sinks until there are none, using the same resampling table. Note that at the end of Algorithm 2, some edges may be marked *unvisited* but are still oriented. Suppose that the output of the completed algorithm is σ' , an orientation of all edges. This completed algorithm is just another implementation of Algorithm 1 with a specific order of resampling bad events. Thus the fact mentioned earlier implies that $\sigma' = \sigma$.

On the other hand, the termination conditions of Algorithm 2 correspond to the cases of Lemma 6. One can show, via a simple induction over the while loop, that at the beginning of the loop, the path P is always a directed path starting from v , and all other visited edges point towards the path P . This implies that Line 13 corresponds to case (b1) in Lemma 6, Line 5 corresponds to case (b2), and exiting the while loop in Line 3 corresponds to case (a). When Algorithm 2 terminates, we have decided whether or not v is a sink. By Lemma 6, this decision stays the same under σ' . As $\sigma' = \sigma$, (2) holds. ◀

We then analyse the efficiency of Algorithm 2. The main bottleneck is when there are degree 2 vertices. It would take $\Omega(\ell^2)$ time to resolve an induced path of length ℓ . We then focus on the case where the minimum vertex degree is at least 3. Note that in this case we have $\Omega_S \neq \emptyset$ for any $S \subseteq V$. For two distributions P and Q over the state space Ω , their total variation distance is defined by $d_{\text{TV}}(P, Q) := \sum_{x \in \Omega} |P(x) - Q(x)|/2$. For two random variables $x \sim P$ and $y \sim Q$, we also write $d_{\text{TV}}(x, y)$ to denote $d_{\text{TV}}(P, Q)$.

We will need the notion of martingales in the next lemma, so let us briefly recap the relevant notions. A sequence of random variables $\{X_i\}_{i \geq 0}$ is a *submartingale* if $\forall n \geq 0$, $\mathbf{E}[|X_n|] < \infty$ and $\mathbf{E}[X_{n+1} \mid X_0, \dots, X_n] \geq X_n$. Submartingales enjoy some nice concentration properties. What we will need is the Azuma-Hoeffding inequality, which states that if $\{X_i\}_{i \geq 0}$ is a submartingale and $|X_{i+1} - X_i| \leq c$ for all $i \geq 0$, then for any $n \geq 0$ and $\varepsilon > 0$,

$$\Pr[X_n - X_0 \leq -\varepsilon] < \exp\left(\frac{-\varepsilon^2}{2nc^2}\right). \quad (3)$$

Then we have the following lemma.

► **Lemma 8** (efficient truncation of Algorithm 2). *Let $G = (V, E)$ be a graph with minimum degree at least 3. Let $S \subseteq V$, $v \notin S$ and $0 < \varepsilon < 1$. Let x be the output of $\text{Sample}(G, S, v)$, and x' constructed as*

- *if $\text{Sample}(G, S, v)$ terminates within $32 \ln(33/\varepsilon)$ executions of Line 12, let $x' = x$;*
- *otherwise, let $x' = 1$.*

Then, it holds that

$$d_{\text{TV}}(x, x') \leq \varepsilon.$$

Proof. We track the length of the path P during the execution of Algorithm 2. When an edge is chosen in Line 4 and sampled in Line 12 of Algorithm 2, the following happens:

- with probability $1/2$, w is appended to P and the length of P increases by one;
- with probability $1/2$, $\{u, w\}$ is marked as **visited**, and the length of P decreases by one in the next iteration if and only if $\{u, w\}$ was the last **unvisited** edge incident to u .

Let X_i be the random variable denoting the length of P after the i -th execution of Line 12 in Algorithm 2. Then the observation above implies that $\{X_i\}_{i \geq 0}$ forms a submartingale. We construct another sequence of random variables $\{Y_i\}_{i \geq 0}$ modified from $\{X_i\}_{i \geq 0}$ as follows:

- $Y_0 = X_0 = 1$.
- At the i -th execution of Line 12 in Algorithm 2:
 - if $\{u, w\}$ is the only **unvisited** edge incident to u , set $Y_{i+1} = X_{i+1} - X_i + Y_i$,
 - otherwise, set $Y_{i+1} = X_{i+1} - X_i + Y_i - 1/2$.

It can be verified that the sequence $\{Y_i\}_{i \geq 0}$ is a martingale.

▷ **Claim 9.** For any $i \geq 0$, $X_i - Y_i \geq i/4$.

Proof. Note that $X_i - Y_i = X_{i-1} - Y_{i-1} + c_i$ where $c_i = 0$ if $\{u, w\}$ is the only **unvisited** edge incident to u at the i -th execution of Line 12 in Algorithm 2, and $c_i = 1/2$ otherwise. Then we can write $X_i - Y_i$ as

$$X_i - Y_i = X_0 - Y_0 + \sum_{j=1}^i c_j.$$

For any i such that $c_i = 0$, let i' be the last index such that $c_{i'} = 0$, or $i' = 0$ if no such i' exists. Since the minimum degree of G is at least 3, when we append any vertex u to P , there are at least two **unvisited** edges incident to u . It implies that there must be some j such that $i' < j < i$ and $c_j = 1/2$. Thus $X_i - Y_i = \sum_{k=1}^i c_k \geq i/4$. ◁

Next we show that if Algorithm 2 doesn't terminate after $32 \ln(33/\varepsilon)$ steps, with high probability the length of the path will not return to 0. As $\{Y_i\}_{i \geq 0}$ is a martingale and $|Y_{i+1} - Y_i| \leq 1$ for all $i \geq 0$, the Azuma–Hoeffding inequality, namely (3), implies that, for any $T > 0$ and $C > 0$,

$$\Pr[Y_T - Y_0 \leq -C] \leq \exp\left(-\frac{C^2}{2T}\right). \quad (4)$$

Thus,

$$\Pr[X_T = 0] \leq \Pr[Y_T \leq -T/4] \leq \Pr[Y_T - Y_0 \leq -T/4] \leq \exp(-T/32),$$

where the first inequality is by Claim 9, and the last inequality is by plugging $C = T/4$ into (4). Then, we have

$$\begin{aligned} \sum_{T=\lceil 32 \ln \frac{33}{\varepsilon} \rceil}^{\infty} \Pr[X_T = 0] &\leq \sum_{T=\lceil 32 \ln \frac{33}{\varepsilon} \rceil}^{\infty} \exp(-T/32) \leq \sum_{T=32 \ln \frac{33}{\varepsilon}}^{\infty} \exp(-T/32) \\ &= \frac{\varepsilon}{33(1 - e^{-1/32})} < \varepsilon. \end{aligned} \quad (5)$$

To finish the proof, we couple x and x' by the same execution of Algorithm 2. Thus, if it terminates within $32 \ln(33/\varepsilon)$ executions of Line 12, then $x = x'$ with probability 1. If not, (5) implies that $x = 0$ with probability at most ε . As we always output $x = 1$ in this case, $x' \neq x$ with probability at most ε , which finishes the proof. ◀

Algorithm 3 SampleEdge(G, S, e).

Input : an undirected graph $G = (V, E)$, a subset of vertices $S \subseteq V$, and an edge $e = \{x, y\} \in E$;

Output : a random orientation $\sigma_e \in \{(x, y), (y, x)\}$ of e ;

- 1 Initialise a mapping $M : E \rightarrow \{\text{visited}, \text{unvisited}\}$ so that $\forall e \in E, M(e) = \text{unvisited}$;
- 2 Let P be a (directed) path initialised as $P = (x, y)$ or $P = (y, x)$ with equal probability, and mark e visited;
- 3 **while** True **do**
 - 4 Let u be the last vertex of P ;
 - 5 **if** $|V(P)| \geq 2$ and $u \notin S$ **then return** the first edge in P ;
 - 6 **if** all edges incident to u are marked visited **then**
 - 7 mark all edges incident to u as unvisited;
 - 8 **if** $|V(P)| = 2$ **then**
 - 9 rerandomise P as $P = (x, y)$ or $P = (y, x)$ with equal probability;
 - 10 **else**
 - 11 remove u from P ;
 - 12 **else**
 - 13 let $e = \{u, w\}$ be the first unvisited edge incident to u ;
 - 14 mark e as visited;
 - 15 **with probability** 1/2 **do**
 - 16 **if** $w \in V(P)$, or there is a visited edge (w, w') for some $w' \in V(P)$ **then**
 - 17 return the first edge in P ;
 - 18 append w to the end of P ;

Note that Lemma 8 does not require $\Omega_S \neq \emptyset$. This is because it is implied by the minimum degree requirement. This implication is an easy consequence of the symmetric Shearer's bound. It is also directly implied by Lemma 10 which we show next.

Lemma 8 implies the first part of Theorem 1. For the second part, we will need a modified version of Algorithm 2 to sample from the marginal distributions of the orientation of edges. This is given in Algorithm 3. It takes as input a subset of vertices $S \subseteq V$ and an edge $e \in E$, then outputs a random orientation σ_e following the marginal distribution induced from μ_S on e . The differences between Algorithm 2 and Algorithm 3 are:

- In Algorithm 2, the number of vertices in P is initialised as $|V(P)| = 1$, while in Algorithm 3, it is initialised as $|V(P)| = 2$.
- Algorithm 2 can terminate when $|V(P)| = 1$ while Algorithm 3 can not terminate due to (a) of Lemma 6 unless x or y are not in S .

The correctness of Algorithm 3 is due to a coupling argument similar to Lemma 7. We couple Algorithm 1 and Algorithm 3 by using the same resampling table. By the same argument as in Lemma 7, given the same resampling table, the orientation of e is the same in the outputs of both Algorithm 1 and Algorithm 3. Thus, σ_e follows the desired marginal distribution by Lemma 5. As for efficiency, we notice that the same martingale argument as in Lemma 8 applies to the length of P as well. Early truncation of the edge sampler only incurs a small error. This finishes the proof of Theorem 1.

4 Applications of the local sampler

We derive the main applications of our local sampler in this section. Lemma 8 implies an additive error on the truncated estimator. As we are after relative errors in approximate counting, we need a lower bound of the marginal ratio.

► **Lemma 10.** *Let $G = (V, E)$ be a graph with a minimum degree at least 3. For any $S \subseteq V$ and $v \notin S$, it holds that $|\Omega_S| \neq 0$ and*

$$\mu_S(v \text{ is not a sink}) > \frac{1}{2}.$$

The proof of Lemma 10 can be viewed as an application of the symmetric Shearer's Lemma [42] on SFO, and is deferred to Section 4.1. Note that the minimum degree requirement is essential for such a marginal lower bound to hold, as the marginal ratio in Lemma 10 can be of order $O(1/n)$ when G is a cycle and $S = V \setminus \{v\}$.

We then show the two approximate counting algorithms first, namely Corollary 2 and Corollary 4.

Proof of Corollary 2. By (1), we just need to $\varepsilon/(2n)$ -approximate $\mu_{V_{i-1}}(v_i \text{ is not a sink})$ for each i to ε -approximate $|\Omega_V|$, the number of sink-free orientations to G . The only random choice Algorithm 2 makes is Line 12. In view of Lemma 8, we enumerate the first $32 \ln(132n/\varepsilon)$ random choices of Algorithm 2, and just output 1 if the algorithm does not terminate by then. Let the estimator be the average of all enumeration. Note that Lemma 10 implies that $\Omega_{V_i} \neq \emptyset$ for any i . Then, Lemmas 7 and 8 imply that the estimator is an $\varepsilon/(4n)$ additive approximation. By Lemma 10, it is also an $\varepsilon/(2n)$ relative approximation, which is what we need.

For the running time, there are n marginals, it takes $\exp(32 \ln(132n/\varepsilon))$ enumerations for each marginal probability, and each enumeration takes at most $O(\ln(132n/\varepsilon))$ time. Therefore, the overall running time is bounded by $O(n(n/\varepsilon)^{32} \log(n/\varepsilon))$. ◀

Proof of Corollary 4. We use (1) again. Denote $\nu_i = \mu_{V_{i-1}}(v_i \text{ is not a sink})$ and $\nu = \prod_{i=1}^n \nu_i$. Let $\tilde{X}_i := \frac{1}{n} \sum_{t=1}^n x'_{i,t}$ be the average of n independent samples from Algorithm 2 truncated after $32 \ln(33 \times 12n/\varepsilon)$ executions of Line 12. Let $\tilde{X} := \prod_{i=1}^n \tilde{X}_i$ be an estimator for ν .

For any i and t , let $\tilde{\nu}_i$ be the expectation of $x'_{i,t}$ (note that it does not depend on t). By Lemmas 7 and 8, $|\tilde{\nu}_i - \nu_i| \leq \frac{\varepsilon}{12n}$. By Lemma 10, $1 - \frac{\varepsilon}{6n} \leq \frac{\tilde{\nu}_i}{\nu_i} \leq 1 + \frac{\varepsilon}{6n}$. Let $\tilde{\nu} = \prod_{i=1}^n \tilde{\nu}_i$ so that $\mathbf{E}[\tilde{X}] = \tilde{\nu}$. Then, as $0 < \varepsilon < 1$,

$$1 - \frac{\varepsilon}{3} \leq \frac{\tilde{\nu}}{\nu} \leq 1 + \frac{\varepsilon}{3}. \quad (6)$$

We bound $\mathbf{Var}[\tilde{X}_i]$ and $\mathbf{Var}[\tilde{X}]$ next. First,

$$\mathbf{Var}[\tilde{X}_i] = \mathbf{Var}\left[\frac{1}{n} \sum_{t=1}^n x'_{i,t}\right] = \frac{1}{n^2} \sum_{t=1}^n \mathbf{Var}[x'_{i,t}] \leq \frac{1}{n},$$

as each $x'_{i,t}$ is an indicator variable. Then,

$$\frac{\mathbf{Var}[\tilde{X}]}{(\mathbf{E}[\tilde{X}])^2} = \frac{\mathbf{E}[\tilde{X}^2]}{(\mathbf{E}[\tilde{X}])^2} - 1 = \frac{\prod_{i=1}^n \mathbf{E}[\tilde{X}_i^2]}{\prod_{i=1}^n (\mathbf{E}[\tilde{X}_i])^2} - 1 = \prod_{i=1}^n \left(1 + \frac{\mathbf{Var}[\tilde{X}_i]}{(\mathbf{E}[\tilde{X}_i])^2}\right) - 1$$

$$\leq \left(1 + \frac{4}{n}\right)^n - 1 < e^4 - 1 < 54. \quad (\text{by Lemma 10})$$

To further reduce the variance, let \hat{X} be the average of N independent samples of \tilde{X} , where $N := \lceil 36 \times 54 / \varepsilon^2 \rceil$. Then, $\mathbf{Var}[\hat{X}] \leq \frac{\mathbf{Var}[\tilde{X}]}{N}$. By Chebyshev's bound, we have

$$\Pr \left[\left| \hat{X} - \tilde{\nu} \right| \geq \frac{\varepsilon}{3} \cdot \tilde{\nu} \right] \leq \frac{9 \mathbf{Var}[\tilde{X}]}{\varepsilon^2 \tilde{\nu}^2} \leq \frac{9 \times 54 \varepsilon^2 \tilde{\nu}^2}{36 \times 54} \cdot \frac{1}{\varepsilon^2 \tilde{\nu}^2} \leq \frac{1}{4}.$$

Thus with probability at least $\frac{3}{4}$, we have that $(1 - \frac{\varepsilon}{3})\tilde{\nu} \leq \hat{X} \leq (1 + \frac{\varepsilon}{3})\tilde{\nu}$. By (6), when this holds, $(1 - \varepsilon)\nu \leq \hat{X} \leq (1 + \varepsilon)\nu$. It is then easy to have an ε -approximation of $|\Omega_V|$.

For the running time, each sample $x'_{i,t}$ takes $O(\log(n/\varepsilon))$ time. We draw n samples for each of the n vertices, and we repeat this process $N = O(\varepsilon^{-2})$ times. Thus, the total running time is bounded by $O((n/\varepsilon)^2 \log(n/\varepsilon))$. ◀

For Corollary 3, we need the second part of Theorem 1 and Algorithm 3. In addition, some extra care for the self-reduction in the overall sampling algorithm is required.

Proof of Corollary 3. We sequentially sample the orientation of edges in G (approximately) from its conditional marginal distribution. Suppose we choose an edge $e = \{u, v\}$, and the sampled orientation is (u, v) . Then, we can remove e from the graph, and let $S \leftarrow S \setminus \{u\}$. The conditional distribution is effectively the same as μ_S in the remaining graph.

One subtlety here is that doing so may create vertices of degree ≤ 2 . To cope with this, we keep sampling edges adjacent to one vertex in S as much as possible before moving on to the next. Suppose the current focus is on v . We use **SampleEdge** to sample the orientation of edges adjacent to v one at a time until either v is removed from S or the degree of v becomes 1. In the latter case, the leftover edge must be oriented away from v , which also results in removing v from S . Note that, when either condition holds, the last edge sampled is oriented as (v, u) for some neighbour u of v . We then move our focus to u if $u \in S$, and move to an arbitrary vertex in S otherwise. The key property of choosing edges this way is that, whenever **SampleEdge**(G, S, e) is invoked, there can only be at most one vertex of degree 2 in S , and if it exists, it must be an endpoint of e . If all vertices are removed from S , we finish by simply outputting a uniformly at random orientation of the remaining edges.

To maintain efficiency, we truncate **SampleEdge**(G, S, e) in each step of the sampling process. More specifically, for some constant C , we output the first edge of P once the number of executions of Line 15 in Algorithm 3 exceeds $C \ln(m/\varepsilon)$. We claim that there is a constant C such that the truncation only incurs an ε/m error in total variation distance between the output and the marginal distribution. This is because the same martingale argument as in Lemma 8 still applies. Note that if P visits any vertex not in S , the algorithm immediately terminates. Thus degrees of vertices not in S do not affect the argument. Moreover, the only degree 2 vertex in S , say x , is adjacent to the first edge $e = \{x, y\}$ of P . If e is initialised as (x, y) , then when P returns to x , the algorithm immediately terminates. Otherwise e is initialised as (y, x) , in which case there is no drift in the first step of P . Thus, as long as we adjust the constant to compensate the potential lack of drift in the first step, the martingale argument in Lemma 8 still works and the claim holds. As the truncation error is ε/m , we may couple the untruncated algorithm with the truncated version, and a union bound implies that the overall error is at most ε .

As we process each edge in at most $O(\log(m/\varepsilon))$ time, the overall running time is then $O(m \log(m/\varepsilon))$. This finishes the proof of the fast sampling algorithm. ◀

4.1 Proof of the marginal lower bound

Now we prove the lower bound of marginal ratios for SFOs, namely, Lemma 10. Let us first recall the variable framework of the local lemma. Consider the probability space \mathcal{P} of a uniformly random orientation of G (namely orienting each edge independently and uniformly at random), and each $u \in S$ corresponding to a bad event \mathcal{E}_u of u being a sink. We then have

$$p_u := \Pr_{\mathcal{P}}[\mathcal{E}_u] = 2^{-d(u)}, \quad \forall u \in V,$$

where $d(u)$ denotes the degree of u . We also need some definitions, essentially from [29] and small variations from those in [42].

► **Definition 11.** We define the following notations.

■ Let $\text{Ind}(G)$ denote all independent sets of G , i.e.,

$$\text{Ind}(G) := \{I \subseteq V \mid I \text{ contains no edge of } G\}.$$

■ For $J \subseteq V$, let

$$q_J := \sum_{\substack{I \in \text{Ind}(G) \\ I \subseteq J}} (-1)^{|I|} \prod_{u \in I} p_u, \quad (7)$$

and

$$P_J := \Pr_{\mathcal{P}} \left[\bigwedge_{u \in J} \neg \mathcal{E}_u \right],$$

which is the probability under \mathcal{P} that all vertices in J are sink-free.

We then proceed to the proof.

Proof of Lemma 10. For $u \in V$, let $\Gamma(u)$ denote the set of neighbours of u in G . We claim that for any $J \subseteq V$ and $u \in J$:

1. $P_J = q_J > 0$;
2. $\frac{q_J}{q_{J \setminus \{u\}}} > \begin{cases} \frac{1}{2} & \Gamma(u) \subseteq J; \\ \frac{2}{3} & \text{otherwise.} \end{cases}$

Lemma 10 immediately follows from the claim as $P_S = \frac{|\Omega_S|}{2^{|E|}} > 0$ and

$$\mu_S(v \text{ is not a sink}) = \Pr_{\mathcal{P}} \left[\neg \mathcal{E}_v \mid \bigwedge_{u \in S} \neg \mathcal{E}_u \right] = \frac{P_{S \cup \{v\}}}{P_S} = \frac{q_{S \cup \{v\}}}{q_S} > \frac{1}{2},$$

where the last equality is by Item 1 and the inequality is by Item 2.

We then prove the claim by induction on the size of J . The base case is when $J = \emptyset$, and all items directly hold as $P_{\emptyset} = q_{\emptyset} = 1$. For the induction step, we first prove Item 1. For $J \subseteq V$ and $u \in J$, denote $\Gamma^+(u) = \Gamma(u) \cup \{u\}$. We have

$$\begin{aligned} \Pr_{\mathcal{P}} \left[\mathcal{E}_u \wedge \left(\bigwedge_{j \in J \setminus \{u\}} \neg \mathcal{E}_j \right) \right] &= \Pr_{\mathcal{P}}[\mathcal{E}_u] \Pr_{\mathcal{P}} \left[\left(\bigwedge_{j \in J \setminus \{u\}} \neg \mathcal{E}_j \right) \mid \mathcal{E}_u \right] \\ &= p_u \Pr_{\mathcal{P}} \left[\left(\bigwedge_{j \in J \setminus \Gamma^+(u)} \neg \mathcal{E}_j \right) \mid \mathcal{E}_u \right] \\ &= p_u \Pr_{\mathcal{P}} \left[\left(\bigwedge_{j \in J \setminus \Gamma^+(u)} \neg \mathcal{E}_j \right) \right] = p_u \cdot P_{J \setminus \Gamma^+(u)}, \end{aligned}$$

where in the second equality we used the fact that S -SFO instances are extremal. Thus,

$$P_J = P_{J \setminus \{u\}} - \Pr_{\mathcal{P}} \left[\mathcal{E}_u \wedge \left(\bigwedge_{j \in J \setminus \{u\}} \neg \mathcal{E}_j \right) \right] = P_{J \setminus \{u\}} - p_u \cdot P_{J \setminus \Gamma^+(u)}. \quad (8)$$

Also, by separating independent sets according to whether they contain u or not, we have

$$\begin{aligned} q_J &= \sum_{\substack{I \in \text{Ind}(G) \\ I \subseteq J}} (-1)^{|I|} \prod_{i \in I} p_i \\ &= \sum_{\substack{I \in \text{Ind}(G) \\ I \subseteq J \setminus \{u\}}} (-1)^{|I|} \prod_{i \in I} p_i - p_u \cdot \sum_{\substack{I \in \text{Ind}(G) \\ I \subseteq J \setminus \Gamma^+(u)}} (-1)^{|I|} \prod_{i \in I} p_i \\ &= q_{J \setminus \{u\}} - p_u \cdot q_{J \setminus \Gamma^+(u)}. \end{aligned} \quad (9)$$

Combining (8), (9), and the induction hypothesis, we have that $P_J = q_J$. For the positivity, let $J \cap \Gamma^+(u)$ be listed as $\{u, u_1, \dots, u_k\}$ for some $k \leq d(u)$. For $0 \leq i \leq k$, let $U_i = \{u, u_1, \dots, u_i\}$ (with $U_0 = \{u\}$). Then we have

$$\frac{q_{J \setminus \{u\}}}{q_{J \setminus \Gamma^+(u)}} = \prod_{i=1}^k \frac{q_{J \setminus U_{i-1}}}{q_{(J \setminus U_{i-1}) \setminus \{u_i\}}} > 2^{-k} \geq 2^{-d(u)} = p_u,$$

where the first inequality is by Item 2 of the induction hypothesis. Thus, $q_J > 0$ by (9) and Item 1 holds.

It remains to show Item 2. Recall that $J \cap \Gamma^+(u)$ is listed as $\{u, u_1, \dots, u_k\}$. For any $0 \leq i \leq k-1$, $\Gamma^+(u_{i+1}) \not\subseteq J \setminus U_i$ because $u \in \Gamma^+(u_{i+1})$ and $u \notin (J \setminus U_i)$. Then by the induction hypothesis on the second case of Item 2,

$$\frac{q_{(J \setminus U_i) \setminus \{u_{i+1}\}}}{q_{J \setminus U_i}} < \frac{3}{2}. \quad (10)$$

By (9), we have

$$\frac{q_J}{q_{J \setminus \{u\}}} = 1 - p_u \cdot \frac{q_{J \setminus \Gamma^+(u)}}{q_{J \setminus \{u\}}} = 1 - 2^{-d(u)} \cdot \prod_{i=0}^{k-1} \frac{q_{(J \setminus U_i) \setminus \{u_{i+1}\}}}{q_{J \setminus U_i}} \stackrel{(10)}{>} 1 - 2^{-d(u)} \cdot \left(\frac{3}{2}\right)^k. \quad (11)$$

If $\Gamma(u) \subseteq J$, $k \leq d(u)$ and as $d(u) \geq 3$,

$$1 - 2^{-d(u)} \cdot \left(\frac{3}{2}\right)^k \geq 1 - \left(\frac{3}{4}\right)^{d(u)} \geq \frac{37}{64} > \frac{1}{2}.$$

If $\Gamma(u) \not\subseteq J$, $k \leq d(u) - 1$ and, again, as $d(u) \geq 3$,

$$1 - 2^{-d(u)} \cdot \left(\frac{3}{2}\right)^k \geq 1 - \frac{2}{3} \cdot \left(\frac{3}{4}\right)^{d(u)} \geq \frac{23}{32} > \frac{2}{3}.$$

Together with (11), this finishes the proof of Item 2 and the lemma. \blacktriangleleft

In the proof above, Item 1 holds mainly because the instance is extremal. For general non-extremal cases, we would have $\frac{P_J}{P_{J \setminus \{u\}}} \geq \frac{q_J}{q_{J \setminus \{u\}}}$ for $J \subseteq V$ and $u \in J$ instead.

4.2 Independence polynomial at negative weights

An interesting consequence of Item 1 in the proof of Lemma 10 is that the number of SFOs can be computed using the independence polynomial evaluated at negative activities. More specifically, similar to (7), let

$$q_G(\mathbf{x}) = \sum_{I \in \text{Ind}(G)} \prod_{u \in I} x_u,$$

where \mathbf{x} is a vector of weights for each vertex. Then, $q_G(-\mathbf{p}) = q_V$ where q_V is defined in (7), and thus $|\Omega_V| = 2^{|E|} q_G(-\mathbf{p})$, where \mathbf{p} is the vector $(p_u)_{u \in V}$ of failure probabilities at the vertices. Namely $p_u = 2^{-d(u)}$ where $d(u)$ is the degree of u .

There are more than one FPTASes [40, 28] that can efficiently approximate the independence polynomial at negative weights. These algorithms work in the so-called Shearer's region [42]. To explain Shearer's region, let us abuse the notation slightly and extend the definition in (7) to a function $q_J(\mathbf{x}) = \sum_{I \in \text{Ind}(G), I \subseteq J} \prod_{u \in I} x_u$ to take an input weight vector \mathbf{x} .

Then, a vector \mathbf{p} is in Shearer's region if and only if $q_S(-\mathbf{p}) > 0$ for all $S \subseteq V$. Lemma 10 implies that the probability vector for SFOs is in Shearer's region. Moreover, we say a vector \mathbf{p} has slack α if $(1 + \alpha)\mathbf{p}$ is in Shearer's region. For a vector \mathbf{x} with slack α , the algorithm by Patel and Regts [40] ε -approximates $q_G(\mathbf{x})$ in time $(n/\varepsilon)^{O(\log d/\alpha)}$, and the algorithm by Harvey, Srivastava, and Vondrák [28] runs in time $(n/(\alpha\varepsilon))^{O(\log d/\sqrt{\alpha})}$, where d is the maximum degree of the graph. They do not recover Corollary 2 as the slack is a constant when constant degree vertices exist. If, in the meantime, some other vertices have unbounded degrees, these algorithms run in quasi-polynomial time instead.

To see the last point, we construct a graph that contains vertices of unbounded degrees but with constant slack for SFOs. Consider the wheel graph G , which consists of a cycle C_n of length n , and a central vertex v that connects to all vertices of C_n . Thus, $p_v = 2^{-n}$ and $p_u = 1/8$ for any u in C_n . For the cycle, as there are two SFOs, we see that $q_{C_n}(-1/4) = 2^{-n+1}$ (where we use Item 1 in the proof of Lemma 10). Thus, by (9),

$$q_G(-2\mathbf{p}) = q_{C_n}(-1/4) - 2p_v = 2^{-n+1} - 2 \cdot 2^{-n} = 0.$$

Therefore, the slack here is at most 1, despite the existence of a high degree vertex.

In summary, the existing FPTASes on the independence polynomial with negative weights do not handle the mixture of high and low degree vertices well enough for the case of SFOs. However, it might provide an alternative approach to derive FPTASes to count solutions to extremal instances of the local lemma, which is worthy of further study.

5 Concluding remarks

Originally, Bubley and Dyer [10] introduced sink-free orientations as a special case of read-twice SAT. Here, "read-twice" means that each variable in a CNF formula appears exactly twice, and it corresponds to an edge of the graph. Vertices of the graph correspond to clauses of the formula. The assignment of the edge is an orientation. This represents that the variable appears with opposite signs in the formula. In fact, Bubley and Dyer showed an FPRAS for all read-twice #SAT. It is natural to ask if they admit FPTAS as well. This question was first raised by Lin, Liu, and Lu [35], who also gave an FPTAS for monotone read-twice #SAT (which is equivalent to counting edge-covers in graphs). The monotone requirement means that the two appearances of any variable have the same sign. From this perspective, our FPTAS is complementary to that of [35]. However, as our techniques are drastically different from [35], to give an FPTAS for all read-twice #SAT, one may need to find a way to combine these two techniques to handle mixed appearances of variables.

Another immediate question is to generalise our local sampler under the partial rejection sampling framework. The first step would be to be able to handle degree 2 vertices for SFOs, which breaks our current submartingale argument. To go a bit further, a local sampler for all extremal instances would yield an FPTAS for all-terminal reliability, whose existence is a major open problem. Also, for all-terminal reliability, one may also attempt to localise the near-linear time sampler in [13].

Lastly, in addition to the discussion of Section 4.2, let us discuss another polynomial associated with SFOs and its zero-freeness. Fix a SFO σ . Let $p(x) = \sum_{i=0}^m C_i x^i$, where m is the number of edges, and C_i indicates how many SFOs exist that agree with σ in exactly i edges. It is easy to evaluate $p(0) = 1$, and $p(1)$ is the total number of SFOs. However, for a cycle, this polynomial becomes $1 + x^m$, which can have a zero arbitrarily close to 1. This zero defeats, at least, the standard application of Barvinok’s method [5, 40]. Although one could exclude cycles by requiring the minimum degree to be at least 3 (like we did in this paper), current techniques of proving zero-freeness seem to hinge on handling all subgraphs. For example, to use Ruelle’s contraction like in [26], one has to start from small fragments of the graph and gradually rebuild it. The obstacle then is to avoid starting from or encountering cycles in the rebuilding process. Other methods, such as the recursion-based one in [36], require hereditary properties (similar to the so-called strong spatial mixing) that break in cycles as well. It would be interesting to see if any of our arguments can help in proving zero-freeness of the polynomial above.

References

- 1 Konrad Anand, Weiming Feng, Graham Freifeld, Heng Guo, and Jiaheng Wang. Approximate counting for spin systems in sub-quadratic time. *TheoretCS*, 4:3:1–3:27, 2025. doi:10.46298/theoretics.25.3.
- 2 Konrad Anand and Mark Jerrum. Perfect sampling in infinite spin systems via strong spatial mixing. *SIAM J. Comput.*, 51(4):1280–1295, 2022. doi:10.1137/21M1437433.
- 3 Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. In *FOCS*, pages 1319–1330. IEEE, 2020. doi:10.1109/FOCS46700.2020.00125.
- 4 Antar Bandyopadhyay and David Gamarnik. Counting without sampling: new algorithms for enumeration problems using statistical physics. In *SODA*, pages 890–899. ACM Press, 2006. URL: <http://dl.acm.org/citation.cfm?id=1109557.1109655>.
- 5 Alexander I. Barvinok. *Combinatorics and Complexity of Partition Functions*, volume 30 of *Algorithms and combinatorics*. Springer, 2016. doi:10.1007/978-3-319-51829-9.
- 6 Mohsen Bayati, David Gamarnik, Dimitriy A. Katz, Chandra Nair, and Prasad Tetali. Simple deterministic approximation algorithms for counting matchings. In *STOC*, pages 122–127. ACM, 2007. doi:10.1145/1250790.1250809.
- 7 Ferenc Bencs and Guus Regts. Barvinok’s interpolation method meets Weitz’s correlation decay approach. *arXiv*, abs/2507.03135, 2025. arXiv:2507.03135.
- 8 Sebastian Brandt, Orr Fischer, Juho Hirvonen, Barbara Keller, Tuomo Lempiäinen, Joel Rybicki, Jukka Suomela, and Jara Uitto. A lower bound for the distributed Lovász local lemma. In *STOC*, pages 479–488. ACM, 2016. doi:10.1145/2897518.2897570.
- 9 Sebastian Brandt, Christoph Grunau, and Václav Rozhon. The randomized local computation complexity of the Lovász local lemma. In *PODC*, pages 307–317. ACM, 2021. doi:10.1145/3465084.3467931.
- 10 Russ Bubley and Martin E. Dyer. Graph orientations with no sink and an approximation for a hard case of #SAT. In *SODA*, pages 248–257. ACM/SIAM, 1997. URL: <http://dl.acm.org/citation.cfm?id=314161.314263>.

- 11 Russ Bubley and Martin E. Dyer. Path coupling: A technique for proving rapid mixing in markov chains. In *FOCS*, pages 223–231. IEEE Computer Society, 1997. doi:10.1109/SFCS.1997.646111.
- 12 Xiaoyu Chen, Weiming Feng, Heng Guo, Xinyuan Zhang, and Zongrui Zou. Deterministic counting from coupling independence. *arXiv*, abs/2410.23225, 2024. doi:10.48550/arXiv.2410.23225.
- 13 Xiaoyu Chen, Heng Guo, Xinyuan Zhang, and Zongrui Zou. Near-linear time samplers for matroid independent sets with applications. In *APPROX/RANDOM*, volume 317 of *LIPICs*, pages 32:1–32:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICs.APPROX/RANDOM.2024.32.
- 14 Zongchen Chen and Yuzhou Gu. Fast sampling of b-matchings and b-edge covers. In *SODA*, pages 4972–4987. SIAM, 2024. doi:10.1137/1.9781611977912.178.
- 15 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Optimal mixing of glauher dynamics: Entropy factorization via high-dimensional expansion. In *STOC*, pages 1537–1550. ACM, 2021. doi:10.1145/3406325.3451035.
- 16 Henry Cohn, Robin Pemantle, and James Gary Propp. Generating a random sink-free orientation in quadratic time. *Electron. J. Comb.*, 9(1), 2002. doi:10.37236/1627.
- 17 Peter Davies-Peck. On the locality of the Lovász local lemma. In *STOC*. ACM, 2025. to appear. doi:10.1145/3717823.3718103.
- 18 Martin E. Dyer, Alan M. Frieze, and Ravi Kannan. A random polynomial time algorithm for approximating the volume of convex bodies. *J. ACM*, 38(1):1–17, 1991. doi:10.1145/102782.102783.
- 19 Paul Erdős and László Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. In *Infinite and finite sets (Colloq., Keszthely, 1973; dedicated to P. Erdős on his 60th birthday)*, Vols. I, II, III, volume Vol. 10 of *Colloq. Math. Soc. János Bolyai*, pages 609–627. North-Holland, Amsterdam-London, 1975.
- 20 Weiming Feng, Heng Guo, Chunyang Wang, Jiaheng Wang, and Yitong Yin. Towards derandomising Markov chain Monte Carlo. In *FOCS*, pages 1963–1990. IEEE, 2023. doi:10.1109/FOCS57990.2023.00120.
- 21 Oded Goldreich, Shafi Goldwasser, and Asaf Nussboim. On the implementation of huge random objects. *SIAM J. Comput.*, 39(7):2761–2822, 2010. doi:10.1137/080722771.
- 22 Heng Guo and Kun He. Tight bounds for popping algorithms. *Random Struct. Algorithms*, 57(2):371–392, 2020. doi:10.1002/rsa.20928.
- 23 Heng Guo and Mark Jerrum. A polynomial-time approximation algorithm for all-terminal network reliability. *SIAM J. Comput.*, 48(3):964–978, 2019. doi:10.1137/18M1201846.
- 24 Heng Guo and Mark Jerrum. Approximately counting bases of bicircular matroids. *Comb. Probab. Comput.*, 30(1):124–135, 2021. doi:10.1017/S0963548320000292.
- 25 Heng Guo, Mark Jerrum, and Jingcheng Liu. Uniform sampling through the Lovász local lemma. *J. ACM*, 66(3):Art. 18, 31, 2019. doi:10.1145/3310131.
- 26 Heng Guo, Chao Liao, Pinyan Lu, and Chihao Zhang. Zeros of holant problems: Locations and algorithms. *ACM Trans. Algorithms*, 17(1):4:1–4:25, 2021. doi:10.1145/3418056.
- 27 Heng Guo and Pinyan Lu. Uniqueness, spatial mixing, and approximation for ferromagnetic 2-spin systems. *ACM Trans. Comput. Theory*, 10(4):17:1–17:25, 2018. doi:10.1145/3265025.
- 28 Nicholas J. A. Harvey, Piyush Srivastava, and Jan Vondrák. Computing the independence polynomial: from the tree threshold down to the roots. In *SODA*, pages 1557–1576. SIAM, 2018. doi:10.1137/1.9781611975031.102.
- 29 Nicholas J. A. Harvey and Jan Vondrák. Short proofs for generalizations of the Lovász local lemma: Shearer’s condition and cluster expansion. *arXiv*, abs/1711.06797, 2017. arXiv:1711.06797.
- 30 Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM J. Comput.*, 18(6):1149–1178, 1989. doi:10.1137/0218077.

- 31 Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM J. Comput.*, 22(5):1087–1116, 1993. doi:10.1137/0222066.
- 32 Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoret. Comput. Sci.*, 43(2-3):169–188, 1986. doi:10.1016/0304-3975(86)90174-X.
- 33 Kashyap Babu Rao Kolipaka and Mario Szegedy. Moser and Tardos meet Lovász. In *STOC*, pages 235–244. ACM, 2011. doi:10.1145/1993636.1993669.
- 34 Liang Li, Pinyan Lu, and Yitong Yin. Correlation decay up to uniqueness in spin systems. In *SODA*, pages 67–84. SIAM, 2013. doi:10.1137/1.9781611973105.5.
- 35 Chengyu Lin, Jingcheng Liu, and Pinyan Lu. A simple FPTAS for counting edge covers. In *SODA*, pages 341–348. SIAM, 2014. doi:10.1137/1.9781611973402.25.
- 36 Jingcheng Liu, Alistair Sinclair, and Piyush Srivastava. A deterministic algorithm for counting colorings with 2-Delta colors. In *FOCS*, pages 1380–1404. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00085.
- 37 Ankur Moitra. Approximate counting, the Lovász local lemma, and inference in graphical models. *J. ACM*, 66(2):10:1–10:25, 2019. doi:10.1145/3268930.
- 38 Robin A. Moser. A constructive proof of the Lovász local lemma. In *STOC*, pages 343–350. ACM, 2009. doi:10.1145/1536414.1536462.
- 39 Robin A. Moser and Gábor Tardos. A constructive proof of the general Lovász local lemma. *J. ACM*, 57(2):11, 2010. doi:10.1145/1667053.1667060.
- 40 Viresh Patel and Guus Regts. Deterministic polynomial-time approximation algorithms for partition functions and graph polynomials. *SIAM J. Comput.*, 46(6):1893–1919, 2017. doi:10.1137/16M1101003.
- 41 Ronitt Rubinfeld, Gil Tamir, Shai Vardi, and Ning Xie. Fast local computation algorithms. In *ICS*, pages 223–238. Tsinghua University Press, 2011. URL: <http://conference.iiis.tsinghua.edu.cn/ICS2011/content/papers/36.html>.
- 42 James B. Shearer. On a problem of Spencer. *Combinatorica*, 5(3):241–245, 1985. doi:10.1007/BF02579368.
- 43 Leslie G. Valiant. The complexity of computing the permanent. *Theor. Comput. Sci.*, 8:189–201, 1979. doi:10.1016/0304-3975(79)90044-6.
- 44 Leslie G. Valiant. The complexity of enumeration and reliability problems. *SIAM J. Comput.*, 8(3):410–421, 1979. doi:10.1137/0208032.
- 45 Dror Weitz. Counting independent sets up to the tree threshold. In *STOC*, pages 140–149. ACM, 2006. doi:10.1145/1132516.1132538.
- 46 David Bruce Wilson. Generating random spanning trees more quickly than the cover time. In *STOC*, pages 296–303. ACM, 1996. doi:10.1145/237814.237880.