

Sampling Partial Acyclic Orientations in Chordal Graphs by the Lovász Local Lemma (Student Abstract)

Wenbo Sun and Ivona Bezáková

Rochester Institute of Technology
102 Lomb Memorial Drive, Rochester, NY 14623
ws3109@rit.edu

Abstract

Sampling of various types of acyclic orientations of chordal graphs plays a central role in several AI applications. In this work we investigate the use of the recently proposed general partial rejection sampling technique of Guo, Jerrum, and Liu, based on the Lovász Local Lemma, for sampling partial acyclic orientations. For a given undirected graph, an acyclic orientation is an assignment of directions to all of its edges so that there is no directed cycle. In partial orientations some edges are allowed to be undirected. We show how the technique can be used to sample partial acyclic orientations of chordal graphs fast and with a clearly specified underlying distribution. This is in contrast to other samplers of various acyclic orientations with running times exponentially dependent on the maximum degree of the graph.

Introduction

Sampling and counting of different types of acyclic orientations of chordal graphs plays an important role in several learning applications, perhaps most notably the structure learning of Bayesian networks that utilizes the so-called v -structure-free acyclic orientations (Ganian, Hamm, and Talvitie 2020; Ghassami et al. 2019; Talvitie and Koivisto 2019). An undirected graph is chordal if every cycle of length four or more has a chord (a non-cycle edge connecting two cycle vertices). A directed graph is v -structure-free if it does not contain a triplet of vertices a, b, c with edges $a \rightarrow c, b \rightarrow c$, and no edge between a and b . To the best of our knowledge, the current state of the art uniform sampling and counting algorithms take time that is exponential in the maximum degree of the underlying chordal graph.

The problems of sampling and counting of acyclic orientations also attracted attention in the theoretical community. Counting of acyclic orientations is a special case of the *Tutte polynomial* of a graph (Welsh 1995). Another special case is the “chromatic polynomial,” which, surprisingly, yields the number of acyclic orientations when evaluated at -1 . While the computation of this polynomial is NP-hard on general graphs, it can be computed efficiently for any chordal graph G , yielding an efficient counting algorithm for acyclic orientations of G . For the so-called self-reducible problems the

existence of an efficient counting algorithm implies an efficient sampler. However, it is unclear if the problem of counting all acyclic orientations is self-reducible, leading to the following open problem: Is it possible to efficiently sample acyclic orientations of chordal graphs uniformly at random?

In this work we investigate what types of acyclic orientations that can be sampled efficiently using the recently proposed general *partial rejections sampling* framework of (Guo, Jerrum, and Liu 2019) based on the Lovász Local Lemma. As a first step in this direction, we study partial acyclic orientations, which are allowed to include directed and undirected edges and contain no directed cycles (Conte et al. 2016). We provide a polynomial time partial-rejection-based sampler for these orientations for distributions that depend on the number of directed edges. Since these orientations can be easily extended into complete acyclic orientations, we hope that this sampler might be used to construct uniform samplers for acyclic orientations and possibly also v -structure free acyclic orientations.

Sampling Partial Acyclic Orientations by Partial Rejection Sampling

The partial rejection sampling framework consists of a set of random variables X_1, \dots, X_n , where each X_i is drawn from its own distribution, and “bad” events A_1, \dots, A_N triggered by specific variable assignments. In its basic form, where simultaneously occurring bad events involve disjoint variables, the partial rejection sampling “rejects” the assignments of the variables in the bad events and re-samples them. If the random variables are drawn from an identical uniform distribution, this process leads to the uniform distribution over assignments with no bad events. One can view the interaction between the bad events through a “dependency graph”: its N vertices correspond to the A_i ’s, and A_i and A_j are connected by an edge if they share one or more variables. Under the *extremal condition* (two bad events are either independent or disjoint), we get a sample drawn according to the product distribution of the individual random variables, scaled to only assignments with no bad events, if we keep re-sampling variables involved in bad events until no bad event occurs. Moreover, this process converges very fast.

(Guo, Jerrum, and Liu 2019) extended the framework beyond the extremal condition, i.e., when two dependent bad

Algorithm 1: Select the resampling set

- 1: Let \mathcal{R} be the set of directed cycles of length three and let \mathcal{N} , the set of events that will not be resampled, be \emptyset . Let $\partial\mathcal{R}$ denote the boundary of \mathcal{R} (neighbors of \mathcal{R} outside \mathcal{R}) in the dependency graph.
 - 2: While $\partial\mathcal{R} \setminus \mathcal{N} \neq \emptyset$, go through each triangle $T \in \partial\mathcal{R} \setminus \mathcal{N}$; if $\text{var}(T) \cap \text{var}(\mathcal{R}) = X_e$ and edge e is directed, then add T into \mathcal{R} . Otherwise add T into \mathcal{N} .
 - 3: Output \mathcal{R} .
-

events can occur at the same time. In addition to resampling the set \mathcal{X} of the variables involved in the current bad events, they also resample the variables for which some values trigger a new bad event with the current assignment of the variables in \mathcal{X} . The resulting distribution is still a product distribution conditioned on no occurring bad events, and the expected running time is stated in the following theorem.

Theorem 0.1 (Guo-Jerrum-Liu). *Let n be the number of variables, N be the number of bad events, and Δ be the maximum degree of the dependency graph. For any bad event, let p be an upper bound on the probability that this bad event occurs. Finally, let r be the maximum probability such that for a pair of neighboring bad events A, B and an assignment of values to the variables in A , if the variables in $\text{var}(B) \setminus \text{var}(A)$ are drawn, B occurs. Then, for any $\Delta \geq 2$, if $6ep\Delta^2 \leq 1$ and $3er\Delta \leq 1$, the expected number of resampling rounds is $O(\log N)$ and thus the expected number of variable resamples is $O(n \log N)$ with high probability.*

Our Contribution

We define a random variable X_e for every edge $e = (u, v)$ of a given chordal graph G . This random variable will take one of three possible values: direction from u to v , direction from v to u , or staying undirected, chosen from distribution $[q, q, 1 - 2q]$ for some $q \in [0, 1]$ that will be specified soon. We define bad events as directed cycles of length three, since any directed cycles of larger sizes can be decomposed because of chordality. We note that the partial acyclic orientations do not fit the extremal condition since two directed cycles of length three can share an edge. Applying the above framework, we keep resampling variables in the output set of Algorithm 1 until no bad event occurs, and the resulting partial acyclic orientations of G is sampled with probability proportional to $q^{\text{number of oriented edges}}$.

Theorem 0.2. *Let $G = (V, E)$ be a chordal graph with maximum degree d . If $q \leq 0.24d^{-\frac{2}{3}}$, the expected number of variable resamples is $O(|E| \log |V|)$.*

Proof. We will apply Theorem 0.1. The probability that a specific bad event occurs is $p = 2q^3$, since the cycle can be oriented in two ways and it has three edges. Let A and B be two neighboring events. This can happen only if they share a single edge. Hence, for B to happen, the shared edge already has to be directed and its other two edges need to take the compatible direction. Thus, $r = q^2$. We claim that $\Delta \leq \frac{3}{2}d$. To see this, let A be a bad event with a

maximum degree in the dependency graph. Let X_{e_1}, X_{e_2} , and X_{e_3} be the variables involved in A and let c_i be the number of other bad events that involve X_{e_i} . These other bad events correspond to different cycles of length three. If $e_1 = (v_1, v_2)$, $e_2 = (v_2, v_3)$, and $e_3 = (v_3, v_1)$, then v_2 's degree is $c_1 + c_2 + 2$, where the $+2$ are for v_1 and v_3 . Since we have $c_1 + c_2 + 2 \leq d$, $c_2 + c_3 + 2 \leq d$, $c_1 + c_3 + 2 \leq d$, and $c_1 + c_2 + c_3 = \Delta$, we get that $2(c_1 + 1 + c_2 + c_3) + 6 = 2\Delta + 6 \leq 3d$, leading to $\Delta \leq \frac{3}{2}d$. Finally, from $6ep\Delta^2 \leq 1$ and $3er\Delta \leq 1$ we get $q \leq \left(3\sqrt[3]{ed\frac{2}{3}}\right)^{-1} \leq 0.24d^{-\frac{2}{3}}$. Since the given chordal graph has $|E|$ edges and at most $\binom{|V|}{3}$ bad events, the expected number of variable resamples is $O(|E| \log |V|)$. \square

Experimental Work and Future Plans

Our work opens many directions for future research: Is it possible to apply the partial rejection sampling framework to directly sample (complete) acyclic orientations? Or v -structure-free orientations? Does sampling of partial orientations and extending them into complete ones yield the uniform distribution? As for our experimental work so far towards (complete) acyclic orientations, instead of resampling directed 3-cycles as indicated by the above analysis, we tried orienting the edges randomly and resampling each edge in a directed cycle. This results in resampling of each strongly connected component, which we repeat until the graph is acyclic. Our preliminary experiments indicate that the process converges fast for chordal graphs. However, this process, while inspired by the LLL framework, does not satisfy all of the LLL conditions and, therefore, the existing theory does not specify the underlying sampling distribution. We plan to study this distribution in the future.

Acknowledgments

This work was supported by the NSF, award DUE-1819546.

References

- Conte, A.; Grossi, R.; Marino, A.; and Rizzi, R. 2016. Listing acyclic orientations of graphs with single and multiple sources. In *LATIN 2016: Theoretical Informatics*, 319–333.
- Ganian, R.; Hamm, T.; and Talvitie, T. 2020. An Efficient Algorithm for Counting Markov Equivalent DAGs. In *Proceedings of AAAI*, 10136–10143.
- Ghassami, A.; Salehkaleybar, S.; Kiyavash, N.; and Zhang, K. 2019. Counting and sampling from Markov equivalent DAGs using clique trees. In *Proc. of AAAI*, 3664–3671.
- Guo, H.; Jerrum, M.; and Liu, J. 2019. Uniform sampling through the Lovasz local lemma. *Journal of the ACM (JACM)* 66(3): 1–31.
- Talvitie, T.; and Koivisto, M. 2019. Counting and sampling Markov equivalent directed acyclic graphs. In *Proceedings of AAAI*, 7984–7991.
- Welsh, D. 1995. Randomised approximation schemes for Tutte-Gröthendieck invariants. In *Discrete Probability and Algorithms*, 133–148. Springer.