58th Annual IEEE Symposium on Foundations of Computer Science

Variable-version Lovász Local Lemma: Beyond Shearer's Bound

Kun He*†, Liang Li‡, Xingwu Liu§†**, Yuyi Wang¶ and Mingji Xia‖†

*CAS Key Lab of Network Data Science and Technology, Institute of Computing Technology, CAS, Beijing, China.
Email: hekun@ict.ac.cn
†University of Chinese Academy of Sciences.
‡Department of Artificial Intelligence, Ant Financial Services Group, Hangzhou, China. Email: liangli.ll@antfin.com
§State Key Laboratory of Computer Architecture, Institute of Computing Technology, CAS, Beijing, China.
Email: liuxingwu@ict.ac.cn**Corresponding author
¶ETH Zürich, Switzerland. Email: yuwang@ethz.ch
‖State Key Laboratory of Computer Science, Institute of Software, CAS, Beijng, China. Email: mingji@ios.ac.cn

***Abstract*—A tight criterion under which the abstract version Lovász Local Lemma (abstract-LLL) holds was given by Shearer [42] decades ago. However, little is known about that of the variable version LLL (variable-LLL) where events are generated by independent random variables, though this model of events is applicable to almost all applications of LLL. We introduce a necessary and sufficient criterion for variable-LLL, in terms of the probabilities of the events and the event-variable graph specifying the dependency among the events. Based on this new criterion, we obtain boundaries for two families of event-variable graphs, namely, cyclic and treelike bigraphs. These are the first two non-trivial cases where the variable-LLL boundary is fully determined. As a byproduct, we also provide a universal constructive method to find a set of events whose union has the maximum probability, given the probability vector and the event-variable graph.**

Though it is #P-hard in general to determine variable-LLL boundaries, we can to some extent decide whether a gap exists between a variable-LLL boundary and the corresponding abstract-LLL boundary. In particular, we show that the gap existence can be decided without solving Shearer's conditions or checking our variable-LLL criterion. Equipped with this powerful theorem, we show that there is no gap if the base graph of the event-variable graph is a tree, while gap appears if the base graph has an induced cycle of length at least 4. The problem is almost completely solved except when the base graph has only 3-cliques, in which case we also get partial solutions.

A set of reduction rules are established that facilitate to infer gap existence of an event-variable graph from known ones. As an application, various event-variable graphs, in particular combinatorial ones, are shown to be gapful/gapless.

I. Introduction

Lovász Local Lemma, or LLL for short, is one of the most important probabilistic methods that has numerous applications since proposed in 1975 by Erdős and Lovász [12]. Basically, LLL aims at finding conditions under which any given set $\mathcal{A}$ of bad events in a probability space can be avoided simultaneously, namely $\mathbb{P}(\cap_{A\in\mathcal{A}}\overline{A}) > 0$. In the most general setting, the dependency among $\mathcal{A}$ is characterized by an undirected graph $G = ([n], E)$, called a dependency graph of $\mathcal{A}$, which satisfies that for any vertex i, A_i is independent of $\{A_j : j \neq i, j \notin \mathcal{N}(i)\}$, where $\mathcal{N}(i)$ stands for the neighborhood of i in G. In this context, finding the conditions on $\mathcal{A}$ is reduced to the fundamental challenge: Given a graph G, determine its abstract interior $\mathcal{I}_a(G)$ which is the set of vectors $\mathbf{p}$ such that $\mathbb{P}\left(\cap_{A\in\mathcal{A}}\overline{A}\right) > 0$ for any event set $\mathcal{A}$ with dependency graph G and probability vector $\mathbf{p}$. *Local* solutions to this problem are collectively called abstract-LLL. The most frequently used abstract-LLL is as follows:

Theorem 1 ([43]): Given a graph $G = ([n], E)$ and a vector $\mathbf{p} \in (0,1)^n$, if there exist real numbers $x_1, ..., x_n \in (0,1)$ such that $p_i \leqslant x_i \prod_{j\in\mathcal{N}(i)}(1 - x_j)$ for any $i \in [n]$, then $\mathbf{p} \in \mathcal{I}_a(G)$.

An exact characterization of $\mathcal{I}_a(G)$ was presented by Shearer [42] over 30 years ago.

Theorem 2 ([42]): Given a graph $G = ([n], E)$ and a vector $\mathbf{p} \in (0,1)^n$, $\mathbf{p} \in \mathcal{I}_a(G)$ if and only if for any $S \in Ind(G)$, $\sum_{T\supseteq S, T\in Ind(G)}(-1)^{|T|-|S|}\prod_{i\in T} p_i > 0$, where $Ind(G)$ is the collection of independent sets of G.

As in Theorem 1 and Theorem 2, only dependency graphs and probabilities of events are involved in abstract-LLL. However, dependency graphs can only capture which events are dependent (more precisely, which events are independent), but not how they are dependent.

A nice model of richer dependency structures is the variable-generated system $\mathcal{A}$ of events, where each event is a constraint on a set $\mathcal{X}$ of independent random variables that can be continuous or discrete. Suppose $\mathcal{A} = \{A_1, ..., A_n\}$ and $\mathcal{X} = \{X_1, ..., X_m\}$. Let $\mathcal{X}_i \subseteq \mathcal{X}$ be a set of variables that completely determines A_i for each $i \in [n]$. The model can be characterized by an event-variable graph which is a bigraph $H = ([n], [m], E)$ where each pair $(i, j) \in [n]\times[m]$ is an edge if and only if $X_j \in \mathcal{X}_i$. Then the fundamental challenge of LLL becomes the VLLL problem as follows: Given a bigraph H, determine its interior $\mathcal{I}(H)$ which is the set of vectors $\mathbf{p}$ such that $\mathbb{P}\left(\cap_{A\in\mathcal{A}}\overline{A}\right) > 0$ for any variable-generated event system $\mathcal{A}$ with event-variable graph H and probability vector $\mathbf{p}$. LLLs solving this problem are

DOI 10.1109/FOCS.2017.48

collectively called variable-LLL.

The model of variable-generated event systems is important, mainly because most applications of LLL have natural underlying independent variables, e.g., hypergraph coloring [29], satisfiability [15], [14], counting solutions to CNF formulas [30], acyclic edge coloring [18], etc. Besides, most results on the algorithmic aspects of LLL are based on this model (see Section I-A). However, there are no special studies on the VLLL problem. A common approach for using LLL in the variable setting is ignoring the variable information and applying abstract-LLL to a dependency graph. This approach only produces results that cannot be better than Shearer's bound. Recently, Harris [22] presents a condition for *lopsided version* [13] of variable-LLL which can go beyond Shearer's criterion, but his condition is based on more information than the event-variable graph (i.e., how events disagree on variables is needed). Thus, the VLLL problem remains open.

Meanwhile, it is widely believed that Shearer's bound is generally not tight for variable-LLL. More precisely, given a bigraph $H = (U, V, E)$, its *base graph* is defined as the graph $G_H = (U, E')$ where two nodes $u_1, u_2 \in U$ are adjacent if and only if u_1, u_2 share some common neighbor in H. A property of base graph is that if H is an event-variable graph of variable-generated event system $\mathcal{A}$, then G_H is a dependency graph of $\mathcal{A}$, which immediately implies that $\mathcal{I}_a(G_H) \subseteq \mathcal{I}(H)$. When $\mathcal{I}_a(G_H) \neq \mathcal{I}(H)$, we say that Shearer's bound is not tight for H, or H has a gap. The only reported bigraph that has a gap is the 4-cyclic one [27], namely a bigraph whose base graph is the 4-cycle. An exact characterization of the conditions for gap existence is far from clear.

Therefore, we try to solve two closely related peoblems:

1) VLLL problem: characterize the interior $\mathcal{I}(H)$ for any bigraph H. Kolipaka et al. [27] have shown that the Moser-Tardos algorithm is efficient up to the Shearer's bound. However, it remains unknown whether the algorithm converges up to the tight bound of variable-LLL and whether it is efficient even beyond Shearer's bound. Moreover, it is widely believed that better bounds can be obtained through variable-LLL for many combinatorial problems, but how much better can it be? A prerequisite for answering these questions is to know what $\mathcal{I}(H)$ is since it tightly upper-bounds the range of variable-LLL.
2) Gap problem: characterize the conditions for a bigraph to have a gap. The status in quo of variable-LLL is to ignore variable information and apply abstract-LLL. This over-simplification generally compromises the power of variable-LLL, but it is lossless and can be safely used when there is no gap. In addition, VLLL problem makes sense only when a gap exists, otherwise it's solved by Shearer's theorem. All this calls for a solution to the gap problem.

A. Related Work

The first result for abstract-LLL was proved by Erdős and Lovász [12] and the first asymmetric one (Theorem 1) was presented in [43]. Though these results are useful, they are not tight in general. A tight, but not local, criterion (Theorem 2) for abstract-LLL was proposed by Shearer [42] over 30 years ago.

Pegden [34][35] introduced *lefthanded-LLL* which does not hold on all dependency graphs, but it is generally tighter than the condition in Theorem 1 and provides a much simpler form of (tight) conditions on special classes of dependency graphs, e.g., chordal graphs. Instead of bounds only working for some dependency graphs, Bissacot et al. [6] proposed to improve Theorem 1 by cluster expansion. Kolipaka [26] further introduced a hierarchy of bounds (e.g., the *clique-LLL*) which can be applied to any dependency graph and are all tighter than the condition in Theorem 1.

Erdös and Spenser [13] introduced lopsided-LLL, which extends the results in [12] to lopsidependency graphs. Scott and Sokal [41] proved that Shearer's condition is tight for lopsided dependency graphs.

There are settings in which Shearer's bound are not tight in general. The best known one may be the variable-generated event systems, whose tight conditions are one of the main contributions of this paper. Harris [22] extended the concept of lopsidependency to variable-LLL, and proposed a condition which can go beyond Shearer's bound in some cases, but not so in general. Note that Harris' bound cannot be applied to standard variable-LLL, because the key concept of *orderability* cannot be defined on event-variable graphs alone.

To make LLL constructive, various sampling algorithms have been proposed so as to avoid all bad events. Algorithm design for LLL is closely related to different bounds mentioned above. Beck [5] first showed that an *algorithmic version LLL (algorithmic-LLL)* is possible and proposed an efficient deterministic sequential algorithm. In that paper, it was required that the degree of the dependency graph under consideration be upper bounded by $2^{n/48}$, which is a very strong restriction. Several work has been done to relax this requirement [11], [32], [37], [38].

Under the model of variable-generated event systems, Moser and Tardos [33] proposed a simple sampling-based algorithm with expected polynomial runtime. Their algorithm is Las Vegas and outputs an assignment to the random variables so as to avoid all bad events. Though a strong model is used, the condition needed in their analysis is the same as Theorem 1 which is even not tight for the abstract-LLL. Pegden [36] proved that Moser and Tardos's algorithm efficiently converges even under the condition of the cluster expansion local lemma. Kolipaka and Szegedy [27] further showed that under the same model, Moser-Tardos algorithm actually works efficiently up to Shearer's bound. Harris [22] presented an algorithm for lopsided version of variable-LLL under the lopsided condition mentioned above. It is still open what conditions are tight for an efficient constructive variable-LLL. Catarata et al. [9] tried experimental methods to observe the possibilities.

Moser-Tardos algorithm can be naturally parallelized because it is not harmful to do sampling for independent events at the same time. Moser and Tardos showed that this parallelization achieves a better expected runtime, but the condition required in their analysis is slightly stronger than that for the sequential case. In fact, parallel algorithms for LLL has been considered much earlier than the invention of Moser-Tardos algorithm [3]. Recently, there are new researches for parallel algorithms inspired by Moser-Tardos algorithm [21], [23]. Besides, algorithmic-LLL has been studied using distributed computation models [7], [10], [16].

Algorithms have also been devised for LLL with dependent variables and other conditions. Harris and Srinivasan [24] first considered the space of permutations. Achlioptas and Iliopoulos [2] studied algorithms specified by certain multigraphs. Frameworks with resampling oracles are also investigated [1], [25], [28].

Actually, variable-LLL has strong connection with sampling. Guo et al. [19] proposed an algorithmic framework, called partial rejection sampling, which establishes this connection in scenarios such as uniform sampling. In a parallel work, Moitra [31] presended an algorithm to approximately sample solutions to general k-CNF under Lovász Local Lemma-like conditions.

Apart from algorithms, LLL has affected (or has been affected by) many other disciplines, in particular physics. For example, alternating-sign independence polynomials of dependency graphs, which is a key element in Shearer's criterion, are also related to the concept of partition functions in statistical physics [40], [20], [44], [45]. Inspired by this connection, cluster expansion local lemma has been proposed [6], and the lower bound of a singularity point in the hard-core lattice gas model has been improved [26]. LLL has also been enriched by the concept of quantum in physics [4], [39], [17].

NOTATION

- $[n]$: the set $\{1, 2, ..., n\}$ for positive integer n.
- $\mathcal{X}, \mathcal{Y}$: sets of mutually independent random variables.
- X, Y: random variables.
- $\mathbf{p}, \mathbf{q}, \mathbf{r}$: vectors of positive real numbers.
- $\phi(\cdot)$: given $\mathbf{p} \in (0, +\infty)^n$, $\phi(\mathbf{p}) \in (0, 1]^n$ is the vector whose i-th entry is $\min\{1, p_i\}$.
- $\mathcal{A}, \mathcal{B}$: sets of events, or sets of cylinders.
- A, B: events, or cylinders.
- $\overline{A}$: the complementary of the event/cylinder A.
- $\mathbb{P}(A)$: the probability of event A.
- $\mathbb{P}(\mathcal{A})$: the vector whose i-th entry is the probability of the i-th event in $\mathcal{A}$.
- μ: Lebesgue measure on Euclidean (sub)spaces.
- $G = (V, E)$: the undirected graph with vertex set V and edge set E.
- $H = (V_1, V_2, E)$: the bigraph with vertex set $V_1 \cup V_2$ and edge set $E \subseteq V_1 \times V_2$. V_1 and V_2 are called the left part and the right part of H, denoted by $L(H)$ and $R(H)$, respectively.
- $\mathcal{N}_G(v)$: the neighborhood of vertex v in graph G, or $\mathcal{N}(v)$ when G is implicit.
- $\mathbb{I}^{\{i\}}$: the unit interval in the i-th dimension of an Euclidean space, or simply $\mathbb{I}$ when i is implicit.
- $\mathbb{I}^S$: the unit cube $\prod_{i \in S} \mathbb{I}^{\{i\}}$, or simply $\mathbb{I}^m$ when $S = [m]$ for some integer m.

II. RESULTS AND DISCUSSION

The main results of this paper are listed and discussed as follows.

Tight condition for variable-LLL: As we mentioned, Shearer's condition is sufficient and necessary for abstract-LLL, but in general it is not tight for variable-LLL. Our first contribution is a sufficient and necessary condition for variable-LLL, namely an exact characterization of $\mathcal{I}(H)$ for any bigraph H. Characterizing $\mathcal{I}(H)$ is equivalent to delimiting its boundary, simply called the boundary of H and denoted by $\partial(H)$, which consists of the vectors $\mathbf{p}$ such that $(1-\epsilon)\mathbf{p} \in \mathcal{I}(H)$ and $(1+\epsilon)\mathbf{p} \notin \mathcal{I}(H)$ for any $\epsilon \in (0, 1)$.

Theorem 3: Given a bigraph $H = ([n], [m], E)$, let $\mathbf{d} = (d_1, ..., d_m)$ where d_j is the degree of the vertex $j \in R(H)$. For any vector $\mathbf{q} \in (0, 1)^n$, $\lambda\mathbf{q}$ lies on the boundary of H if and only if λ is the optimal solution to the program:

$$
\begin{aligned}
\min \quad & \lambda \\
\text{s.t.} \quad & \sum_{i \in [n]} C_{i,k_1,k_2,...,k_m} \geqslant 1 \text{ for any } k_j \in [d_j], j \in [m]; \\
& \text{For any } (i, j) \in ([n] \times [m]) \backslash E, C_{i,k_1,k_2,...,k_m} \text{ does} \\
& \text{ not depend on } k_j; \\
& \sum_{k_1 \in [d_1],...,k_m \in [d_m]} (\prod_{j \in [m]} x_{jk_j}) C_{i,k_1,k_2,...,k_m} = \lambda q_i \\
& \text{ for } i \in [n]; \\
& \sum_{k \in [d_j]} x_{jk} = 1 \text{ for } j \in [m]; \\
& x_{jk} \in [0, 1] \text{ for } j \in [m], k \in [d_j]; \\
& C_{i,k_1,k_2,...,k_m} \in \{0, 1\} \text{ for } i \in [n], k_j \in [d_j], j \in [m].
\end{aligned}
$$

As far as we know, this is the first condition for general variable-LLL. It essentially means that the variables can be discretized. Namely, to determine the boundary vectors, it is enough to consider the discrete variables taking d_j values. In addition, the program facilitates to construct the "worst-case" set of events, which means that the probability of the union of the events is maximized.

Boundary of cyclic bigraphs: Though the program above is hard to solve in general, its insight of discretization makes it possible to fully determine the boundary of any cyclic bigraph as in the following theorem. Here a bigraph is called n-cyclic if its base graph is a cycle of length n.

Theorem 4: Given a vector $\mathbf{p} \in (0,1)^n$, for each $i \in [n]$, let λ_i be the minimum positive solution to the equation system: $b_1 = \lambda p_i, b_k = \frac{\lambda p_{k+i-1}}{1-b_{k-1}}$ for $2 \leqslant k \leqslant n-1$, $b_{n-1} = 1 - \lambda p_{i-1}$. Let $\lambda_0 = \min_{i\in[n]} \lambda_i$. Then $\lambda_0 \mathbf{p}$ lies on the boundary of any n-cyclic bigraph.

In the literature, only one vector on the boundary of 4-cyclic bigraphs has been identified. The above theorem shows that the whole boundary of any n-cyclic bigraph can be determined by solving an $(n-1)$-degree polynomial equation.

Not only for cyclic bigraphs, we also give a procedure to exactly determine the boundary of treelike bigraphs. A bigraph is called treelike if its base graph is a tree.

A sufficient and necessary condition for gap existence: Since a bigraph provides more information than its base graph, it is naturally expected to have a gap, namely Shearer's bound is not tight for birgpahs. We propose a necessary and sufficient condition to decide whether such a gap exist. For conciseness of presentation, we also call a bigraph gapful if it has a gap, and gapless otherwise.

Theorem 5: Given a bigraph H and a vector $\mathbf{p}$ of positive reals, the following three conditions are equivalent:

1) For any λ such that $\lambda\mathbf{p} \in \mathcal{I}(H)$, there is an exclusive variable-generated event system $\mathcal{A}$ with event-variable graph H and probability vector $\lambda\mathbf{p}$.
2) For the λ such that $\lambda\mathbf{p} \in \partial(H)$, there is an exclusive variable-generated event system $\mathcal{A}$ with event-variable graph H and probability vector $\lambda\mathbf{p}$.
3) H is gapless in the direction of $\mathbf{p}$.

Here the qualifier "exclusive" means that the events in $\mathcal{A}$ are either independent or disjoint, and "gapless in the direction of $\mathbf{p}$" means that for any λ, $\lambda\mathbf{p} \in \mathcal{I}(H)$ if and only if $\lambda\mathbf{p} \in \mathcal{I}_a(G_H)$.

By this criterion, one can check the existence of a gap just by examining the bigraph, without computing Shearer's bound of its base graph.

On this basis, we investigate gap existence for two families of bigraphs.

Theorem 6: Treelike bigraphs are gapless.

Based on this theorem, we develop a simple algorithm to efficiently compute Shearer's bound for any dependency graph which is a tree.

In contrast, we obtain an opposite result for cyclic bigraphs, which considerably extends the only gap-existing example in literature [27].

Theorem 7: Cyclic bigraphs are gapful.

Reduction method: To discover more instances that have or have no gaps, we propose a set of reduction rules which allow us transforming a bigraph without changing the existence or nonexisence of a gap. We identify five basic operations. Three of them as well as their inverses preserve both gapful and gapless; the other two preserve gapful, while the inverses of the two preserve gapless. Applying these operations, we show that a bigraph is gapful if it contains a gapful one. This, together with Theorem 7, intuitively means that Shearer's criterion is not tight for almost all cases of variable-LLL. Likewise, we show that combinatorial bigraphs $H_{n,m}$ are gapful if m is small enough and are gapless if m is large enough.

III. Probability Boundary of Variable-LLL

This section aims at solving the *VLLL problem*: given a bigraph H, determine all the vectors $\mathbf{p}$ such that $\mathbb{P}\left(\cap_{A\in\mathcal{A}}\overline{A}\right) > 0$ for any variable-generated event system $\mathcal{A}$ with event-variable graph H and probability vector $\mathbf{p}$. Basically, we will transform the problem into a geometric one and solve it in the framework of Euclidean geometry.

For conciseness of presentation, a variable-generated event system $\mathcal{A}$ is said to conform with a bigraph H, denoted by $\mathcal{A} \sim H$, if H is an event-variable graph of $\mathcal{A}$.

Throughout this section, we only consider bigraphs whose base graphs are connected. This restriction does not lose generality for the following reason. If a bigraph has disconnected base graph, itself must also be disconnected and each component is again a bigraph. In this case, the interior of the original bigraph is exactly the direct product of the interiors of the component bigraphs.

A. A Geometric Counterpart

Now we formulate a geometric counterpart of the VLLL problem, called the GLLL problem. Consider the m-dimensional Euclidean space $\mathbb{R}^m$ endowed with Lebesgue measure μ. Let X_i be the coordinate variable of the i-th dimension, $i \in [m]$. For any $S \subseteq [m]$, the *S-unit cube*, denoted by $\mathbb{I}^S$, is defined to be the $|S|$-dimensional unit hypercube $[0,1]^{|S|}$ working as the domain range of the variables $\{X_i : i \in S\}$ such that for each $i \in S$, $X_i \in [0,1]$. When $S = [k]$ for some $k \leqslant m$, we simply write $\mathbb{I}^k$ for $\mathbb{I}^{[k]}$. A *cylinder* A in $\mathbb{I}^m$ is a subset of the form $B \times \mathbb{I}^S$, where $B \subseteq \mathbb{I}^{[m]\setminus S}$ is called a base of A; define $\dim(B) = [m]\setminus S$. Given a bigraph $H = ([n],[m],E)$ and a set $\mathcal{A}$ of cylinders $A_1,...,A_n$ in $\mathbb{I}^m$, we say that $\mathcal{A}$ *conforms with* H, also denoted by $\mathcal{A} \sim H$, if there are bases $B_1,...,B_n$ of $A_1,...,A_n$ such that $E = \{(i,j) \in [n]\times[m] : j \in \dim(B_i)\}$. Now comes the **GLLL problem**: given bigraph H, determine all the vectors $\mathbf{p}$ such that $\mu(\cup_{A\in\mathcal{A}}A) < 1$ for any cylinder set $\mathcal{A} \sim H$ with $\mu(\mathcal{A}) = \mathbf{p}$.

One can easily see that the VLLL problem is equivalent to the GLLL problem in the sense that they have the same solutions. Hence, the rest of the paper will be presented in the context of the GLLL problem. To ease understanding, the terms "event" and "cylinder" will be used interchangeably, and so will "probability" and "Lebesgue measure". The complementary of a cylinder A in $\mathbb{I}^{[m]}$ is defined to be the cylinder $\overline{A} = \mathbb{I}^{[m]} \backslash A$.

B. A Sufficient and Necessary Criterion

Definition 1 (Interior): The *interior* of a bigraph H, denoted by $\mathcal{I}(H)$, is the set of vectors $\mathbf{p}$ on $(0,1)$ such that $\mu\left(\cap_{A\in\mathcal{A}}\overline{A}\right) > 0$ for any cylinder set $\mathcal{A} \sim H$ with $\mu(\mathcal{A}) = \mathbf{p}$.

Definition 2 (Exterior): The *exterior* of a bigraph $H = ([n],[m],E)$, denoted by $\mathcal{E}(H)$, is the set $(0,1]^n \backslash \mathcal{I}(H)$.

Definition 3 (Boundary): The *boundary* of a bigraph H, denoted by $\partial(H)$, is the set of vectors $\mathbf{p}$ on $(0,1]$ such that $(1-\epsilon)\mathbf{p} \in \mathcal{I}(H)$ and $(1+\epsilon)\mathbf{p} \notin \mathcal{I}(H)$ for any $\epsilon \in (0,1)$. Any $\mathbf{p} \in \partial(H)$ is called a *boundary vector* of H.

We can show that there is a boundary vector in every direction.

Lemma 8: Given a bigraph $H = ([n],[m],E)$, for any $\mathbf{p} \in (0,1]^n$, there exists a unique $\lambda > 0$ such that $\lambda\mathbf{p} \in \partial(H)$.

In the rest of this section, we propose a program to characterize boundary vectors. The cornerstone of the program is the observation that cylinders can be properly *discretized* without changing the boundary.

Given an integer $d > 0$, a cylinder $A \subseteq \mathbb{I}^m$ is said to be d-discrete in dimension j, if there is a partition of $\mathbb{I}^{\{j\}}$ into d disjoint intervals $\Delta_1,...,\Delta_d$ such that $A = \cup_{k=1}^{d} S_k^A \times \Delta_k$ for some $S_k^A \subseteq \mathbb{I}^{[m]\backslash\{j\}}$, $k = 1,\ldots,d$. A cylinder set $\mathcal{A}$ is called d-discrete in dimension j, or discrete in dimension j when d is implicit, if so is every $A \in \mathcal{A}$. Given a vector $\mathbf{d} = (d_1,...,d_m)$, a cylinder A is called $\mathbf{d}$-discrete, if it is d_j-discrete in dimension j for any $j \in [m]$. Likewise, $\mathcal{A}$ is called $\mathbf{d}$-discrete, or discrete when $\mathbf{d}$ is implicit, if so is every $A \in \mathcal{A}$; then the vector $\mathbf{d}$ is called a discreteness degree of $\mathcal{A}$.

Given two vectors $\mathbf{p}$ and $\mathbf{q}$, we say $\mathbf{p} \leqslant \mathbf{q}$ if the inequality holds entry-wise. Additionally, if the inequality is strict on at least one entry, we say that $\mathbf{p} < \mathbf{q}$.

In the rest of this section, fix a bigraph $H = ([n],[m],E)$ and a probability vector $\mathbf{p} \in \partial(H)$. Let $\mathbf{q}_\epsilon \triangleq \phi((1+\epsilon)\mathbf{p})$ for any real number $\epsilon > 0$ and $\mathbf{d} \triangleq (d_1,...,d_m)$ with each d_j being the degree of the vertex $j \in [m]$ in H.

The main results (Theorem 13 and Theorem 3) of this section present a discrete cylinder set for each probability vector on the boundary. As a byproduct, it is shown that the boundary lies in the exterior. Following these theorems, there are two corollaries handling the discretization of interior and exterior respectively.

The boundary is discretized in four steps, as shown in the coming four lemmas. First, we show that for any $\epsilon > 0$, there is a discrete cylinder set whose measure vector lies in the exterior and is ϵ-close to $\mathbf{p}$. Unfortunately, the discreteness degree of this cylinder set depends on ϵ, and may be unbounded when ϵ tends to 0. Second, we show that the set of cylinders can be chosen such that the discreteness degree is no more than $\mathbf{d}$. However, the measure vector may not be lower-bounded by $\mathbf{p}$, though it is still upper-bounded by $\mathbf{q}_\epsilon$. Third, with ϵ tending to 0, a mathematic program and a calculus argument guarantee the existence of a $\mathbf{d}$-discrete cylinder set whose measure vector lies in the exterior and is upper-bounded by $\mathbf{p}$. Finally, we show that the measure vector of this cylinder set is exactly $\mathbf{p}$, which immediately leads to the main theorem.

The basic idea of proving the next lemma is to discretize cylinders dimension by dimension. To discretize the j-th dimension, the axis $\mathbb{I}^{\{j\}}$ is partitioned so that every cylinder varies little in each part, which naturally leads to an approximation (that is discrete in dimension j) to the origin cylinders. The partition is found by approximating an integral with a finite summation.

Lemma 9: For any $\epsilon > 0$, there exists a discrete cylinder set $\mathcal{A} \sim H$ such that $\mathbf{p} \leqslant \mu(\mathcal{A}) \leqslant \mathbf{q}_\epsilon$ and $\mu(\cup_{A\in\mathcal{A}}A) = 1$.

The basic idea of proving the next lemma is as follows. By Lemma 9, we have a discrete cylinder set. The vector of the measures of the cylinders that depend on a common variable X_j turns out to be a convex combination of d_j-dimensional vectors. A simple combinatorial argument indicates that at most d_j out of the latter vectors are enough to generate (also by convex combination) former one, which immediately implies the desired discreteness degree.

Lemma 10: For any $\epsilon > 0$, there exists a $\mathbf{d}$-discrete cylinder set $\mathcal{A} \sim H$ such that $\mu(\mathcal{A}) \leqslant \mathbf{q}_\epsilon$ and $\mu(\cup_{A\in\mathcal{A}}A) = 1$.

By Lemma 10, for any small $\epsilon > 0$, there is a $\mathbf{d}$-discrete cylinder set $\mathcal{A}_\epsilon$ whose measure is upper bounded by $\mathbf{q}_\epsilon$. The next lemma claims that this is the case even if $\epsilon = 0$. The basic idea is to show that as ϵ tends to 0, $\mathcal{A}_\epsilon$ converges in some sense and the limit is a $\mathbf{d}$-discrete cylinder set. For this end, we establish an equivalence between the existence of a $\mathbf{d}$-discrete cylinder set and a mathematical program consisting of polynomial constraints. This equivalence, together with an argument based on the continuity of the constraints, ensures that a sequence of $\mathcal{A}_\epsilon$ converges and the limit cylinder set is as desired.

Lemma 11: There is a $\mathbf{d}$-discrete cylinder set $\mathcal{A} \sim H$ such that $\mu(\mathcal{A}) \leqslant \mathbf{p}$ and $\mu(\cup_{A\in\mathcal{A}}A) = 1$.

For the cylinder set $\mathcal{A}$ obtained in Lemma 11, the next lemma claims that $\mu(\mathcal{A}) = \mathbf{p}$. Roughly speaking, if there are A_i and A_j both depending on X_l and satisfying that $\mu(A_i) < p_i$ and $\mu(A_j) = p_j$, we can remove a thin slice (perpendicular to the axis X_l) from A_j and attach it to A_i. After this operation, both $\mu(A_i) < p_i$ and $\mu(A_j) < p_j$, no extra dependency is brought about, and the whole cube remains been filled up. Iteratively, we can finally get $\mu(A_k) < p_k$ for any k, which is contradictory to the assumption that $\mathbf{p}$ is a boundary vector.

Lemma 12: If there is a cylinder set $\mathcal{A} \sim H$ such that $\mu(\mathcal{A}) \leqslant \mathbf{p}$ and $\mu(\cup_{A\in\mathcal{A}}A) = 1$, then $\mu(\mathcal{A}) = \mathbf{p}$.

Our main theorem immediately follows from Lemma 11 and Lemma 12.

Theorem 13: Given a bigraph $H = ([n], [m], E)$ and $\mathbf{p} \in \partial(H)$, let $\mathbf{d} = (d_1, ..., d_m)$ where d_j is the degree of the vertex $j \in R(H)$. Then there is a $\mathbf{d}$-discrete cylinder set $\mathcal{A} \sim H$ such that $\mu(\mathcal{A}) = \mathbf{p}$ and $\mu(\cup_{A \in \mathcal{A}} A) = 1$.

Theorem 13 and Lemma 12 essentially give a **necessary and sufficient condition for deciding the boundary**: $\mathbf{p}$ is a boundary vector if and only if it is a minimal probability vector that allows a cylinder set as in Theorem 13. Due to discreteness, such cylinders have only finitely many forms, so their existence can be checked at least by the exhaustive method. In this sense, not only can we decide boundary vectors, but also **constructively** find the "worst-case" cylinders (i.e., the measure of the union is maximized). The method is as in Theorem 3.

Theorem 3: Given a bigraph $H = ([n], [m], E)$, let $\mathbf{d} = (d_1, ..., d_m)$ where d_j is the degree of the vertex $j \in R(H)$. For any vector $\mathbf{q} \in (0, 1)^n$, $\lambda\mathbf{q}$ lies on the boundary of H if and only if λ is the optimal solution to the program:

$$\begin{aligned} \min \quad & \lambda \\ \text{s.t.} \quad & \sum_{i \in [n]} C_{i,k_1,k_2,...,k_m} \geqslant 1 \text{ for any } k_j \in [d_j], j \in [m]; \\ & \text{For any } (i, j) \in ([n] \times [m]) \backslash E, C_{i,k_1,k_2,...,k_m} \text{ does} \\ & \text{ not depend on } k_j; \\ & \sum_{k_1 \in [d_1],...,k_m \in [d_m]} \Big(\prod_{j \in [m]} x_{jk_j}\Big) C_{i,k_1,k_2,...,k_m} = \lambda q_i \\ & \text{ for } i \in [n]; \\ & \sum_{k \in [d_j]} x_{jk} = 1 \text{ for } j \in [m]; \\ & x_{jk} \in [0, 1] \text{ for } j \in [m], k \in [d_j]; \\ & C_{i,k_1,k_2,...,k_m} \in \{0, 1\} \text{ for } i \in [n], k_j \in [d_j], j \in [m]. \end{aligned}$$

Given a solution to the program, $\mathbb{I}^m$ is partitioned into subcubes by cutting every axis X_j into d_j intervals of length x_{jk_j}, $k_j \in [d_j]$. For each $i \in [n]$, let A_i be the union of the subcubes numbered by $(k_1, k_2, ..., k_m)$ with $C_{i,k_1,k_2,...,k_m} = 1$. Then $\mathcal{A} = \{A_1, ..., A_n\}$ satisfies the requirement of Theorem 13.

By Theorem 13, for $\mathbf{p} \in \partial(H)$, the worst set of cylinders can be $\mathbf{d}$-discrete. We will generalize the result to non-boundary vectors. When $\mathbf{p}$ is in the interior of H, the basic idea of the next corollary is to add an extra cylinder to the original set of cylinders so that their union has measure 1. By minimizing the extra cylinder, the union of the original cylinders should be maximized. Then the discreteness degree follows from Theorem 13.

Corollary 14: Given a bigraph $H = ([n], [m], E)$ and $\mathbf{p} \in \mathcal{I}(H)$, define $\mathbf{d} = (d_1, ..., d_m)$ where d_j is the degree of the vertex $j \in R(H)$. Let $\mathbf{d}' = (d_1 + 1, ..., d_m + 1)$. Then there is a $\mathbf{d}'$-discrete cylinder set $\mathcal{B} = \arg\max_{\mathcal{A} \sim H, \mu(\mathcal{A}) = \mathbf{p}} \mu(\cup_{A \in \mathcal{A}} A)$.

The next corollary indicates that for $\mathbf{p} \in \mathcal{E}(H)$, the discreteness degree is also small. The basic idea is opposite to that of proving Corollary 14. Some events and/or a part of one are *removed* so that the remaining events exactly fill the cube. Then discretize the rest events according to Theorem 13. Finally, a slight refinement of the discretization also discretizes the removed events.

Corollary 15: Given a bigraph $H = ([n], [m], E)$ and $\mathbf{p} \in \mathcal{E}(H)$, define $\mathbf{d} = (d_1, ..., d_m)$ where d_j is the degree of the vertex $j \in R(H)$. There is a $\tilde{\mathbf{d}}$-discrete cylinder set $\mathcal{A} \sim H$ such that $\mu(\mathcal{A}) = \mathbf{p}$ and $\mu(\cup_{A \in \mathcal{A}} A) = 1$, where $\tilde{d}_{j_0} = d_{j_0} + 1$ for some $j_0 \in [m]$ and $\tilde{d}_j = d_j$ for $j \neq j_0$.

Remark 1: The above theorems and corollaries mean that given a bigraph and a vector in $(0, 1]^n$, the worst case cylinders can be discretized. More importantly, the discreteness degree is determined by the bigraph only.

The discreteness degrees mentioned in Theorems 13, 3 and Corollaries 14, 15 are tight in general. For example, consider the complete bigraph $H = ([n], [1], E)$. For any $\mathbf{p} \in (0, 1]^n$, $\mathbf{p} \in \mathcal{I}(H)$ if and only if $\sum_{i \in [n]} p_i < 1$, while $\mathbf{p} \in \partial(H)$ if and only if $\sum_{i \in [n]} p_i = 1$. One can easily check that the discreteness degrees in the Theorems and Lemmas are the smallest possible for this example.

IV. Breaking cycles

In this section, we compute the boundary of cyclic bigraphs. Roughly speaking, a cyclic bigraph models the variable-generated system of events where events are located on a cycle and neighbors (and only neighbors) depend on common variables. Note that the only gapful bigraph reported in the literature is 4-cyclic [27].

Definition 4 (Cyclic bigraph): A bigraph H is said to be n-cyclic if the base graph G_H is a cycle of length n. When $n = 3$, it additionally requires $\cap_{i \in L(H)} \mathcal{N}_H(i) = \varnothing$. In case of no ambiguity, an n-cyclic bigraph is simply called a cyclic bigraph.

As far as the GLLL problem is concerned, an n-cyclic bigraph is always equivalent to the canonical one $H_n = ([n], [n], E_n)$ where $E_n = \{(i, i), (i, (i + 1)(\overline{\text{mod}}\ n)) : i \in [n]\}$. Here the value $k(\overline{\text{mod}}\ n)$ is defined to be $(k - 1)(\text{mod}\ n) + 1$. Hence, in the rest of this section, we will focus on H_n.

To simplify notation, the operator "$(\overline{\text{mod}}\ n)$" will be omitted whenever clear from context.

A concept that is opposite to cyclic bigraphs is as follows.

Definition 5 (Linear bigraph): A bigraph H is said to be n-linear if the base graph G_H is a path of length n. In case of no ambiguity, an n-linear bigraph is simply called linear.

A rather surprising phenomenon of cyclic bigraphs is that they can be reduced to linear bigraphs in the following sense: any boundary vector of a cyclic bigraph is also that of a linear one. That is, to find the boundary vector in a certain direction, some pair of neighboring events can be decoupled (i.e., get independent of each other) by ignoring their shared variables. In this sense we say that the cycle is broken. The result is stated in the next theorem.

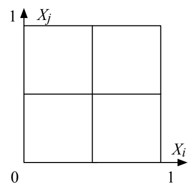

Figure 1. Partitioning $\mathbb{I}^{\{i,j\}}$ into four rectangles

Theorem 16: For any vector $\mathbf{p} \in \partial(H_n)$, there is a $\mathbf{d}$-discrete cylinder set $\mathcal{A} \sim H_n$ such that $\mu(\mathcal{A}) = \mathbf{p}$, $\mu(\cup_{A \in \mathcal{A}} A) = 1$, and $\mathbf{d} < (2, 2, ..., 2)$.

Remark 2: $\mathbf{d} < (2, 2, ..., 2)$ means that $d_j = 1$ for some $j \in [n]$. Then all the cylinders (especially A_j and A_{j+1}) are independent of X_j. As a result, $\mathcal{A}$ also conforms with $H_n^{\{j\}}$, the n-linear bigraph obtained by removing the vertex $j \in R(H_n)$, meaning that $\mathbf{p} \in \mathcal{E}(H_n^{\{j\}})$. Due to the assumption that $\mathbf{p} \in \partial(H_n)$ and the easy fact that $\mathcal{E}(H_n^{\{j\}}) \subseteq \mathcal{E}(H_n)$, $\mathbf{p}$ must also lie on the boundary of $H_n^{\{j\}}$.

Now we give a sketchy proof of Theorem 16.

Arbitrarily fix $\mathbf{p} \in \partial(H_n)$. By Theorem 13, there is a $(2, 2, ..., 2)$-discrete cylinder set $\mathcal{A} \sim H_n$ such that $\mu(\mathcal{A}) = \mathbf{p}$ and $\mu(\cup_{A \in \mathcal{A}} A) = 1$. Arbitrarily choose such a cylinder set $\mathcal{A} = \{A_1, ..., A_n\}$. For each $i \in [n]$, let B_i be the base of A_i such that $\dim(B_i) = \{i, i+1\}$.

Due to the discreteness of $\mathcal{A}$, each $\mathbb{I}^{\{i,j\}}$ is partitioned into four rectangles as in Figure 1 and only unions of some of the rectangles make sense. Especially interesting is the 14 types of non-trivial unions, namely T_{11} through to T_{44}, grouped into the four categories $T_1, ..., T_4$, as shown in Figure 2.

For any $i \in [n]$, since $0 < \mu(B_i) < 1$, B_i must have one of the 14 types in $\mathbb{I}^{\{i,i+1\}}$. Now, we explore how B_i, B_{i+1} are correlated in terms of their types.

If some B_i has type T_2, then A_i is independent of either X_i or X_{i+1}. It is easy to see that $\mathcal{A}$ is 1-discrete either in dimension i or in dimension $i+1$. Hence we have

Lemma 17: If B_i has type T_2 for some $i \in [n]$, $\mathcal{A}$ has a discreteness degree smaller than $(2, 2, ..., 2)$.

As a result, in the rest of this section, it is assumed that no bases have type T_2. Then we show that there are at most two essentially different possibilities of the types of $B_1, ..., B_n$, as indicated in the next lemma.

Lemma 18: There are at most two possible combinations of the types of the bases.

1) T_1-dominant: all bases have type T_1 except one has type T_4.
2) T_3-dominant: all bases have type T_3.

However, the two possibilities are ruled out by the following lemma.

Lemma 19: Neither T_1-dominant nor T_3-dominant combination is possible.

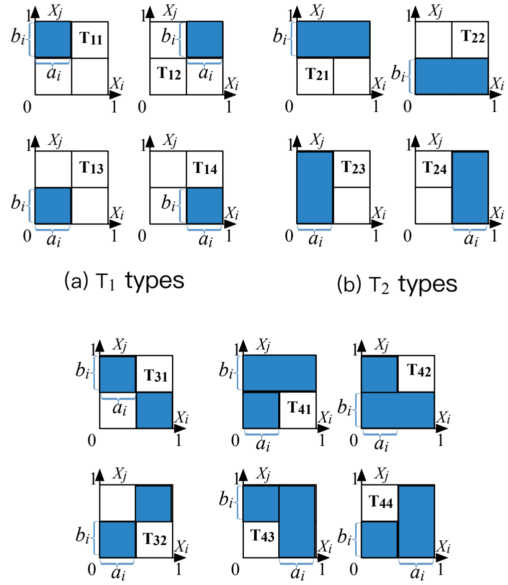

(a) T_1 types (b) T_2 types

(c) T_3 types (d) T_4 types

Figure 2. Types T_{11} through to T_{44} in $\mathbb{I}^{\{i,j\}}$, indicated by the shaded areas

By Lemmas 17 through to 19, Theorem 16 immediately follows, which in turn enables us to find out all boundary vectors of any cyclic bigraph.

Theorem 4: Given a vector $\mathbf{p} \in (0,1)^n$, for each $i \in [n]$, let λ_i be the minimum positive solution to the equation system: $b_1 = \lambda p_i, b_k = \frac{\lambda p_{k+i-1}}{1-b_{k-1}}$ for $2 \leqslant k \leqslant n-1$, $b_{n-1} = 1 - \lambda p_{i-1}$. Let $\lambda_0 = \min_{i \in [n]} \lambda_i$. Then $\lambda_0 \mathbf{p}$ lies on the boundary of any n-cyclic bigraph.

As an application of Theorem 4, we explicitly characterize the boundary of the 3-cyclic bigraph H_3.

Example 1: For H_3, consider an arbitrary $\mathbf{p} \in (0,1)^3$ with $p_1 + p_2 + p_3 = 1$. For $i \in \{1, 2, 3\}$, we have $\lambda_i = \frac{1-\sqrt{1-4p_ip_{i-1}}}{2p_ip_{i-1}}$. Since the function $\frac{1-\sqrt{1-2x}}{x}$ is increasing with $x > 0$, the final λ_0 is the λ_i with i minimizing p_ip_{i-1}. For example, if $p_1 \geqslant p_2$ and $p_1 \geqslant p_3$, then $\lambda_3 \mathbf{p} = \frac{1-\sqrt{1-4p_2p_3}}{2p_2p_3}\mathbf{p}$ is a boundary vector.

V. Gap between Abstract- and Variable-LLL

In this section, we investigate conditions under which Shearer's bound remains tight for Variable-LLL.

A. A Theorem for Gap Decision

Definition 6 (Exclusiveness): An event set $\mathcal{A}$ is said to be *exclusive* with respect to a graph G, if G is a dependency graph of $\mathcal{A}$ and $\mu(A_i \cap A_j) = 0$ for any i, j such that $i \in \mathcal{N}_G(j)$. A cylinder set $\mathcal{A}$ is called exclusive with respect to a bigraph H, if $\mathcal{A}$ conforms with H and $\mathcal{A}$ is exclusive with respect to G_H. We do not mention "with respect to G or H" if it is clear from context.

The next lemma claims that exclusive cylinder sets always exist if only the measures are small enough.

Lemma 20: For any bigraph H, there is $\epsilon > 0$ such that for any vector $\mathbf{p}$ on $(0, \epsilon)$, there exists a cylinder set $\mathcal{A}$ that is exclusive with respect to H and $\mu(\mathcal{A}) = \mathbf{p}$.

Definition 7 (Abstract Interior): The *abstract interior* of a graph $G = ([n], E)$, denoted by $\mathcal{I}_a(G)$, is the set $\{\mathbf{p} \in (0,1)^n : \mathbb{P}\left(\cap_{A \in \mathcal{A}} \overline{A}\right) > 0$ for any event set $\mathcal{A} \sim_a G$ with $\mathbb{P}(\mathcal{A}) = \mathbf{p}\}$, where "$\mathcal{A} \sim_a G$" means G is a dependency graph of $\mathcal{A}$. Given a bigraph H, we simply write $\mathcal{I}_a(H)$ for $\mathcal{I}_a(G_H)$.

It is obvious that $\mathcal{I}_a(H) \subseteq \mathcal{I}(H)$ for any bigraph H.

Definition 8 (Abstract Boundary): The *abstract boundary* of a graph $G = ([n], E)$, denoted by $\partial_a(G)$, is the set $\{\mathbf{p} \in (0,1]^n : (1-\epsilon)\mathbf{p} \in \mathcal{I}_a(G)$ and $(1+\epsilon)\mathbf{p} \notin \mathcal{I}_a(G)$ for any $\epsilon \in (0,1)\}$. Any $\mathbf{p} \in \partial_a(G)$ is called an *abstract boundary vector* of G.

The following is a counterpart of Lemma 8.

Lemma 21: For any graph $G = ([n], E)$ and $\mathbf{p} \in (0,1)^n$, there is a unique $\lambda > 0$ such that $\lambda \mathbf{p} \in \partial_a(G)$.

The proof of [42, Theorem 1] presents an interesting property of exclusive event sets as in Lemma 22.

Lemma 22: Given a graph G and $\mathbf{p} \in \mathcal{I}_a(G) \cup \partial_a(G)$. Among all event sets $\mathcal{A} \sim_a G$ with $\mathbb{P}(\mathcal{A}) = \mathbf{p}$, there is an exclusive one such that $\mathbb{P}(\cup_{A \in \mathcal{A}} A)$ is maximized.

Definition 9 (Gap): A bigraph H is called *gapful in the direction* of $\mathbf{p} \in (0,1)^n$, if there is $\lambda > 0$ such that $\lambda \mathbf{p} \in \mathcal{I}(H) \backslash \mathcal{I}_a(H)$, otherwise it is called gapless in this direction. H is said to be *gapful* if it is gapful in some direction, otherwise it is gapless.

For convenience, "being gapful" will be used interchangeably with "having a gap".

The main result of this section, namely Theorem 5, is a necessary and sufficient condition for deciding whether a bigraph is gapful. Intuitively, it bridges gaplessness and exclusiveness both in interior and on boundary. At the first glance, the connection between gaplessness and exclusiveness seems to be an immediate corollary of the well-known Lemma 22 by Shearer. However, this is not the case. The main difficulty lies in boundary vectors. Suppose the bigraph is gapless. On the one hand, for a vector on its boundary, there is an exclusive event set whose union has probability 1, by Lemma 22. These events are not necessarily cylinders, so we can't claim the existence of an exclusive cylinder set. On the other hand, there indeed is a cylinder set whose union has measure 1. Such a cylinder set must be exclusive as desired, if the union of non-exclusive events always has smaller probability than that of exclusive ones. But Lemma 22 just claims that the union of non-exclusive events can't have bigger probability, not precluding the possibility that the probabilities are equal. Our proof essentially distills down to ruling out this possibility, as in Lemma 23.

The following lemma is key to the proof of Theorem 5. Intuitively, it claims that the overall probability is maximized by and only by an exclusive set of event. Note that the "by" part was proved in [42, Theorem 1], but it apears here to make this paper self-contained.

Lemma 23: Suppose that G is a dependency graph of event sets $\mathcal{A}$ and $\mathcal{B}$, $\mathbb{P}(\mathcal{A}) = \mathbb{P}(\mathcal{B})$, and $\mathcal{B}$ is exclusive. Then $\mathbb{P}(\cup_{A \in \mathcal{A}} A) \leqslant \mathbb{P}(\cup_{B \in \mathcal{B}} B)$, and the equality holds if and only if $\mathcal{A}$ is exclusive.

Now we are ready to prove the main theorem of this section.

Theorem 5: Given a bigraph H and a vector $\mathbf{p}$ of positive reals, the following three conditions are equivalent:

1) For any λ such that $\lambda \mathbf{p} \in \mathcal{I}(H)$, there is an exclusive variable-generated event system $\mathcal{A}$ with event-variable graph H and probability vector $\lambda \mathbf{p}$.
2) For the λ such that $\lambda \mathbf{p} \in \partial(H)$, there is an exclusive variable-generated event system $\mathcal{A}$ with event-variable graph H and probability vector $\lambda \mathbf{p}$.
3) H is gapless in the direction of $\mathbf{p}$.

Proof: ($1 \Rightarrow 3$): Arbitrarily fix $\lambda > 0$ such that $\mathbf{q} \triangleq \lambda \mathbf{p} \in \mathcal{I}(H)$. Let $\mathcal{A} \sim H$ be an exclusive cylinder set such that $\mu(\mathcal{A}) = \mathbf{q}$ and $\mu(\cup_{A \in \mathcal{A}} A) < 1$. It also holds that $\mathcal{A}$ is exclusive with respect to the base graph G_H. Since $\mu(\cup_{A \in \mathcal{A}} A) < 1$, by Lemma 23, $\mu(\cup_{B \in \mathcal{B}} B) < 1$ for any event set $\mathcal{B} \sim_a G_H$ with $\mathbb{P}(\mathcal{B}) = \mathbf{q}$. As a result, $\mathbf{q} \in \mathcal{I}_a(H)$. Altogether, H is gapless in the direction of $\mathbf{p}$.

($3 \Rightarrow 2$): Assume that H is gapless in the direction of $\mathbf{p}$. Let λ be such that $\mathbf{q} \triangleq \lambda \mathbf{p} \in \partial(H)$. By Theorem 13, there is a cylinder set $\mathcal{A} \sim H$ such that $\mu(\mathcal{A}) = \mathbf{q}$ and $\mu(\cup_{A \in \mathcal{A}} A) = 1$. On the other hand, $\mathbf{q} \in \partial_a(H)$ due to the assumption that H is gapless in the direction of $\mathbf{p}$. By Lemma 22, there is an exclusive event set $\mathcal{B} \sim_a G_H$ such that $\mu(\mathcal{B}) = \mathbf{q}$ and $\mathbb{P}(\cup_{B \in \mathcal{B}} B) = 1$. Because $\mathcal{A}$ also conforms with G_H and $\mathbb{P}(\cup_{B \in \mathcal{B}} B) = \mathbb{P}(\cup_{A \in \mathcal{A}} A) = 1$, by Lemma 23, $\mathcal{A}$ must be exclusive with respect to G_H, hence exclusive with respect to H.

($2 \Rightarrow 1$): Arbitrarily fix $\lambda > 0$ such that $\mathbf{q} \triangleq \lambda \mathbf{p} \in \mathcal{I}(H)$. Let $\delta > 1$ be such that $\delta \lambda \mathbf{p} \in \partial(H)$. Arbitrarily choose an exclusive cylinder set $\mathcal{A} \sim H$ which satisfies $\mu(\mathcal{A}) = \delta \lambda \mathbf{p}$. Let $\mathcal{A} = \{A_1, ..., A_n\}$. For each $i \in L(H)$, there is a base B_i of A_i such that $\dim(B_i) = \mathcal{N}_H(i)$. Arbitrarily choose a subset $B_i' \subset B_i$ with $\mu(B_i') = \mu(B_i)/\delta$. Let $\mathcal{A}' = \{A_1', ..., A_n'\}$ where each A_i' is the cylinder with base B_i'. It is easy to check that $\mathcal{A}' \sim H$, $\mu(\mathcal{A}') = \mathbf{q}$, and $\mathcal{A}'$ is exclusive. ∎

The significance of Theorem 5 lies in that it enables to decide whether a gap exists without checking Shearer's bound.

Remark 3: Given a bigraph $H = ([n], [m], E)$ and a vector $\mathbf{p} \in (0,1)^n$, consider three real numbers that are of special interest. λ_1, λ_2 are such that $\lambda_1 \mathbf{p} \in \partial(H)$ and $\lambda_2 \mathbf{p} \in \partial_a(G_H)$, respectively. λ_3 is the maximum λ such that there is an exclusive cylinder set $\mathcal{A} \sim H$ with $\mu(\mathcal{A}) = \lambda \mathbf{p}$. It is not difficult to see that $\lambda_1 \geqslant \lambda_2 \geqslant \lambda_3$. An equivalent form of Theorem 5 is that the three numbers are either all equal or pairwise different.

B. Reduction Rules

Given a bigraph H, we define the following 5 types of operations on H.

1) Delete-Variable: Delete a vertex $j \in R(H)$ with $|\mathcal{N}(j)| \leqslant 1$, and remove the incident edge if any.
2) Duplicate-Event: Given a vertex $i \in L(H)$, add a vertex i' to $L(H)$, and add edges incident to i' so that $\mathcal{N}(i') = \mathcal{N}(i)$.
3) Duplicate-Variable: Given a vertex $j \in R(H)$, add a vertex j' to $R(H)$, and add some edges incident to j' so that $\mathcal{N}(j') \subseteq \mathcal{N}(j)$.
4) Delete-Edge: Delete an edge from E provided that the base graph remains unchanged.
5) Delete-Event: Delete a vertex $i \in L(H)$, and remove all the incident edges.

We also define the inverses of the above operations. The inverse of an operation O is the operation O' such that for any H, $O'(O(H)) = O(O'(H)) = H$.

The next theorems show how these operations influence the existence of gaps.

Theorem 24: A bigraph H is gapful, if and only if it is gapful after applying Delete-Variable, Duplicate-Event, Duplicate-Variable, or their inverse operations.

Theorem 25: A gapless bigraph remains gapless after applying Delete-Event or the inverse of Delete-Edge.

Theorem 26: A gapful bigraph remains gapful after applying Delete-Edge or the inverse of Delete-Event.

Because the operations can be pipelined, applying them in combination may produce interesting results. The following corollaries are some examples.

Definition 10 (Combinatorial bigraph): Given two positive integers $m < n$, let $H_{n,m} = ([\binom{n}{m}], [n], E_{n,m})$ where $(i, j) \in E_{n,m}$ if and only if j is in the m-sized subset of $[n]$ represented by i. $H_{n,m}$ is called the (n, m)-combinatorial bigraph.

Corollary 27: If $H_{n,m}$ is gapless, then so is $H_{n+c,m+c}$ for any integer $c \geqslant 1$.

Corollary 28: If $H_{n,m}$ is gapful, then for any integer $c \geqslant 1$, $H_{cn,cm}$ is also gapful.

Definition 11 (Sparsified bigraphs): A bigraph $H' = ([n'], [m'], E')$ is called a sparsification of $H = ([n], [m], E)$ if $[n'] = [n], [m'] \subseteq [m], E' \subseteq E$ and their base graphs are the same.

By Theorem 24 and Theorem 25, we know that if H is gapful, all sparsifications of H must be gapful. Applying Corollary 28, we get the following result.

Corollary 29: If $H_{n,m}$ is gapful, all sparsifications of $H_{cn,cm}$ are also gapful for any integer $c \geqslant 1$.

VI. Relationship between gaps and cycles

In this section, we show that a bigraph has a gap is almost equivalent to that its base graph has an cycle. The only case that is not completely known is when the bigraph does not contain any cyclic bigraph but its base graph has a 3-clique. Many examples in this case is gapless, but we find one that turns out to be gapful.

First of all, we prove that any treelike bigraph is gapless. Recall that a bigraph is called treelike if its base graph is a tree. Basically, for a vector on boundary, we construct an exclusive cylinder set, which leads to the result by Theorem 5. To ensure exclusiveness, the unit interval in each dimension is divided into two disjoint parts, each of which is assigned to one of the two cylinders depending on this dimension. The construction is feasible because the base graph is a tree.

Theorem 6: Treelike bigraphs are gapless.

Using of the constructed cylinders, we get a system of equations whose solution determines the boundary of a treelike bigraph.

Corollary 30: Given a bigraph $H = ([n], [m], E)$ such that G_H is a tree, appoint the vertex n as the root of G_H. For any $\mathbf{p} \in (0, 1)^n$, $\lambda \mathbf{p} \in \partial(H)$ if and only if λ is the minimum positive solution to the equation system: $q_i = \lambda p_i$ if vertex i is a leaf of G_H, $q_i = \lambda p_i / \prod_{k \text{ is a child of } i}(1 - q_k)$ if $i \neq n$ and is not a leaf, and $\lambda p_n = \prod_{k \text{ is a child of } n}(1 - q_k)$.

Now we show that cyclic bigraphs are gapful. Though in principle this can be shown by a combination of [42, Theorem 1] and the results in Section IV, it is tough since both Shearer's inequality system and the high degree polynomial in Theorem 4 are hard to solve. Hence we do it in another way. Specifically, for the vector $\mathbf{q} = (\frac{1}{4} + \epsilon, ..., \frac{1}{4} + \epsilon)$ where $\epsilon > 0$ is small enough, we show two facts. First, the vector $\mathbf{q}$ lies in the interior of the cyclic bigraph. Second, $\mathbf{q}$ does not allow any exclusive cylinder set. By Theorem 5, these facts immediately imply Theorem 7.

Theorem 7: Cyclic bigraphs are gapful.

By Theorem 7, we get a large class of gapful bigraphs.

Definition 12 (Containing): We say that a bigraph H contains another bigraph H', if there are injections $\pi_L : L(H') \to L(H)$ and $\pi_R : R(H') \to R(H)$ such that the following two conditions hold simultaneously:

1) For any $i \in L(H')$ and $j \in R(H')$, $\pi_R(j) \in \mathcal{N}_H(\pi_L(i))$ if and only if $j \in \mathcal{N}_{H'}(i)$.
2) For any $j \in R(H) \backslash \pi_R(R(H'))$, $j \notin \mathcal{N}_H(\pi_L(i)) \cap \mathcal{N}_H(\pi_L(k))$ for any $i, k \in L(H')$.

Intuively, H contains H' means that H' can be *embedded* in H without incurring extra dependency.

By Theorem 24 and Theorem 25, a bigraph H is gapful if it contains a gapful one. According to Theorem 7, we obtain the following result.

Corollary 31: Any bigraph containing a cyclic one is gapful.

Based on Theorem 6 and Corollary 31, it is natural to have the following conjecture:

Conjecture 1 (Gap conjecture): A bigraph is gapful if and only if it contains a cyclic bigraph.

We have already known that the sufficiency does hold. As to the necessity, assume that the bigraph H does not contain any cyclic one. We analyze case by case.

Case 1: the base graph is a tree. By Theorem 6, H is gapless, as desired.

Case 2: the base graph has cycles. Since H does not contain a cyclic bigraph, its base graph does not have induced cycles longer than three. As a result, solving the conjecture is equivalent to answering the following question Q**: Is a bigraph gapless if it does not contain any cyclic one but its base graph has 3-cliques?**

First look at a trivial example: the bigraph $H = ([3], [1], E)$ with $E = [3] \times [1]$. It satisfies the condition of question Q. One can easily check that $\partial(H) = \{(p_1, p_2, p_3) : p_1 + p_2 + p_3 = 1\} = \partial_a(G_H)$. So, H is gapless.

For more evidence, recall $H_{n,m}$, the (n, m)-combinatorial bigraph. As a special case, $H_{3,2}$ is the canonical 3-cyclic bigraph H_3. Generally, we have the following observations:

First, $m = 1$: Only sets of independent events can conform with $H_{n,m}$.

Second, $2 \leqslant m \leqslant \frac{2}{3}n$: $H_{n,m}$ contains 3-cyclic bigraphs, so it is gapful.

Third, $m > \frac{2}{3}n$: $H_{n,m}$ does not contain cyclic bigraphs, but the base graph have 3-cliques since it is a complete graph. We mainly consider bigraphs in this category.

Theorem 32: $H_{4,3}$ is gapless.

Theorem 32, together with Corollary 27, immediately implies the following result.

Corollary 33: For $n \geqslant 4$, $H_{n,n-1}$ is gapless.

Actually, Corollary 33 can be generalized to $H_{n,n-m}$ for any fixed m and large enough n, as shown in Theorem 34.

Theorem 34 is proved by construction. Basically, given a boundary vector $\mathbf{p}$ of $H_{n,n-m}$, we identify a small number of dimensions, partition the unit cube $\mathcal{C}$ spanned by these dimensions into $\binom{n}{m}$ parts, and use each part as the base to construct a cylinder in $\mathbb{I}^n$. Essentially this means projecting all cylinders to a low-dimensional cube. For this end, we first show that when n is big enough, there are 10 dimensions such that any cylinder independent of at least one of these dimensions has very small probability. Then Lemma 20 ensures that the bases of these cylinders can be chosen as exclusive. Finally, the other cylinders are obtained by partitioning the part of $\mathcal{C}$ that has not yet been covered. Altogether, we get an exclusive set of cylinders whose measure vector is $\mathbf{p}$.

Theorem 34: For any constant m, when n is large enough, $H_{n,n-m}$ is gapless.

In spite of so much confirmative evidence, the general answer to the question Q turns out to be NO! The following bigraph is an example where gap is not caused by containing cyclic bigraphs. Specifically, it is the bigraph $H^* = ([5], [5], E)$ with $E = (\{1\} \times \{1, 4, 5\}) \cup (\{2\} \times \{2, 4, 5\}) \cup (\{3\} \times \{3, 4, 5\}) \cup (\{4\} \times \{1, 2, 3, 4\}) \cup (\{5\} \times \{1, 2, 3, 5\})$.

Theorem 35: H^* is gapful.

By Delete-Event and the inverse operation of Duplicate-Variable, it is not difficult to reduce $H_{7,5}$ to H^*. Because H^* is gapful, $H_{7,5}$ is also gapful. From Corollary 29, we have the following corollary.

Corollary 36: For any integer $c \geqslant 1$, every sparsification of $H_{7c,5c}$ is gapful.

In summary, we get some instances of gapful/gapless bigraphs, listed in Table I.

Table I
INSTANCES OF GAPFUL/GAPLESS BIGRAPHS

Gapful	Gapless
H^* Sparsifications of $H_{7c,5c}$ cyclic bigraphs	$H_{n,n-c}$ for large n $H_{n,n-1}$ for $n \geqslant 4$ treelike bigraphs

VII. Hardness Results

We define some computational problems that are closely related to the variable-LLL problem and show that they are difficult to solve.

Definition 13 (MUP Problem): Given a bigraph $H = ([n], [m], E)$ and vector $\mathbf{p} \in (0, 1]^n$, compute $\Psi(H, \mathbf{p}) \triangleq \max_{\mathcal{A} \sim H, \mu(\mathcal{A}) = \mathbf{p}} \mu(\cup_{A \in \mathcal{A}} A)$, where $\mathcal{A}$ ranges on sets of cylinders in $\mathbb{I}^m$ and μ is Lebesgue measure.

Definition 14 (INT Problem): Given a bigraph H and a vector $\mathbf{p}$ on $(0, 1)$, decide whether $\mathbf{p} \in \mathcal{I}(H)$.

Theorem 37: MUP is #P-hard.

Proof: It is enough to show that MUP is #P-hard even if H is a $(3, 2)$-regular bigraph and $\mathbf{p} = (\frac{1}{8}, \frac{1}{8}, ..., \frac{1}{8})$.

Arbitrarily fix a $(3, 2)$-regular bigraph $H = ([n], [m], E)$. We will construct a set $\mathcal{A}$ of cylinders in $\mathbb{I}^m$ such that $\mu(\mathcal{A}) = \mathbf{p}$, $\mathcal{A}$ is exclusive with respect to H, and the probability of the union is maximized.

Arbitrarily choose a function $f : [m] \to [n]$ which maps each vertex in $[m]$ to one of its neighbors. For each $i \in [n]$, a cylinder $A_i \subset \mathbb{I}^m$ is defined in this way: for each neighbor j of i, $0 \leqslant X_j \leqslant 1/2$ if $f(j) = i$, otherwise $1/2 < X_j \leqslant 1$. Let $\mathcal{A} = \{A_1, ..., A_n\}$.

Since each $i \in [n]$ has exactly three neighbors in H and each $j \in [m]$ has exactly two neighbors, we observe that $\mu(\mathcal{A}) = \mathbf{p}$ and $\mathcal{A}$ is exclusive with respect to H. Hence $\mu(\cup_{i \in [n]} A_i) = \Psi(H, \mathbf{p})$, by Lemma 23.

The construction actually partitions $\mathbb{I}^m$ into 2^m blocks each having measure 2^{-m}. Any cylinder in $\mathcal{A}$ consists of some blocks. Let $B_{k_1,k_2,...,k_m}$, $k_j \in \{0, 1\}$ for any j, denote the block defined by $0 \leqslant X_j \leqslant 1/2$ if $k_j = 0$ or $1/2 < X_j \leqslant 1$ if $k_j = 1$, for any $j \in [m]$. Given $k_1, k_2, ..., k_m \in \{0, 1\}$ and $i \in [n]$, the following two conditions are equivalent.

1. $B_{k_1,k_2,...,k_m} \subseteq A_i$.

2. For each neighbor j of i in H, $k_j = 0$ if and only if $f(j) = i$.

Let N be the number of blocks outside of $\cup_{i\in[n]} A_i$. Then we have $\mu(\cup_{i\in[n]} A_i) = 1 - N/2^m$, so computing $\Psi(H, \mathbf{p})$ is equivalent to computing N.

On the other hand, computing N is related to the 3SAT problem. Let $\{y_1, ..., y_m\}$ be a set of boolean variables. For each $i \in [n]$, assume j_1, j_2, j_3 are its neighbors in H; define a 3SAT clause $\phi_i = z_{j_1} \vee z_{j_2} \vee z_{j_3}$ where the literal $z_{j_k} \triangleq y_{j_k}$ if $f(j_k) = i$, otherwise $z_{j_k} \triangleq \overline{y_{j_k}}$, for $k = 1, 2, 3$. The constraint-variable graph of $\phi \triangleq \phi_1 \wedge ... \wedge \phi_n$ is H. Note that each variable appears twice oppositely, so ϕ is a Holant$([0, 1, 0]|[0, 1, 1, 1])$ or Rtw-Opp-#3SAT instance.

Now consider an assignment $y_j = k_j, j \in [m]$. It is straightforward to check that ϕ is satisfied if and only if the block $B_{k_1,k_2,...,k_m}$ is outside $\cup_{i\in[n]} A_i$. Thus, N is the number of satisfying assignments of ϕ, which is #P-hard to compute even if H is (3,2)-regular, by [8, Theorem 8.1]. ■

By Theorem 37, one can prove the following result.

Theorem 38: INT is #P-hard.

ACKNOWLEDGEMENT

We are grateful to the anonymous referees for detailed corrections and suggestions. We also thank Prof. Zhiwei Xu, Prof. Yungang Bao, Prof. Xiaoming Sun and Prof. Roger Wattenhofer for their encouragement and support.

Kun He's work is supported by the National Key Research and Development Program of China (Grant No. 2016YFB1000200) and the National Natural Science Foundation of China (Grant No. 61433014, 61502449, 61602440) and the China National Program for support of Top-notch Young Professionals. Xingwu Liu's work is partially supported by the National Natural Science Foundation of China (Grant No. 61420106013). Yuyi Wang's work is partially supported by the National Natural Science Foundation of China (Grant no. 11601375). Mingji Xia's work is supported by China National 973 Program (Grant No. 2014CB340300).

REFERENCES

[1] D. Achlioptas and F. Iliopoulos. Focused stochastic local search and the Lovász local lemma. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2024–2038. Society for Industrial and Applied Mathematics, 2016.

[2] D. Achlioptas and F. Iliopoulos. Random walks that find perfect objects and the Lovász local lemma. *Journal of the ACM (JACM)*, 63(3):22, 2016.

[3] N. Alon. A parallel algorithmic version of the local lemma. *Random Structures & Algorithms*, 2(4):367–378, 1991.

[4] A. Ambainis, J. Kempe, and O. Sattath. A quantum Lovász local lemma. *Journal of the ACM (JACM)*, 59(5):24, 2012.

[5] J. Beck. An algorithmic approach to the Lovász local lemma. *Random Structures & Algorithms*, 2(4):343–365, 1991.

[6] R. Bissacot, R. Fernández, A. Procacci, and B. Scoppola. An improvement of the Lovász local lemma via cluster expansion. *Combinatorics, Probability and Computing*, 20(05):709–719, 2011.

[7] S. Brandt, O. Fischer, J. Hirvonen, B. Keller, T. Lempiinen, J. Rybicki, J. Suomela, and J. Uitto. A lower bound for the distributed Lovász local lemma. In *48th Annual Symposium on Theory of Computing (STOC), Cambridge, Massachusetts, USA*, June 2016.

[8] J. Cai, P. Lu, and M. Xia. Holographic algorithms by fibonacci gates and holographic reductions for hardness. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 644–653. IEEE Computer Society, 2008.

[9] J. D. Catarata, S. Corbett, H. Stern, M. Szegedy, T. Vyskocil, and Z. Zhang. The Moser-Tardos resample algorithm: Where is the limit?(an experimental inquiry). In *2017 Proceedings of the Ninteenth Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 159–171. SIAM, 2017.

[10] K.-M. Chung, S. Pettie, and H.-H. Su. Distributed algorithms for the Lovász local lemma and graph coloring. In *Proceedings of the 2014 ACM symposium on Principles of distributed computing*, pages 134–143. ACM, 2014.

[11] A. Czumaj and C. Scheideler. A new algorithm approach to the general Lovász local lemma with applications to scheduling and satisfiability problems. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 38–47. ACM, 2000.

[12] P. Erdős and L. Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. *Infinite and finite sets*, 10(2):609–627, 1975.

[13] P. Erdős and J. Spencer. Lopsided Lovász local lemma and latin transversals. *Discrete Applied Mathematics*, 30(2-3):151–154, 1991.

[14] H. Gebauer, R. A. Moser, D. Scheder, and E. Welzl. The Lovász local lemma and satisfiability. In *Efficient Algorithms*, pages 30–54. Springer, 2009.

[15] H. Gebauer, T. Szabó, and G. Tardos. The local lemma is asymptotically tight for SAT. *Journal of the ACM (JACM)*, 63(5):43, 2016.

[16] M. Ghaffari. An improved distributed algorithm for maximal independent set. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 270–277. Society for Industrial and Applied Mathematics, 2016.

[17] A. Gilyén and O. Sattath. On preparing ground states of gapped hamiltonians: An efficient quantum Lovász local lemma. In *Proceedings of the 58th Annual Symposium on Foundations of Computer Science (FOCS 2017)*.

[18] I. Giotis, L. Kirousis, K. I. Psaromiligkos, and D. M. Thilikos. Acyclic edge coloring through the Lovász local lemma. *Theoretical Computer Science*, 665:40–50, 2017.

[19] H. Guo, M. Jerrum, and J. Liu. Uniform sampling through the Lovász local lemma. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 342–355. ACM, 2017.

[20] A. Guttmann. Comment on'the exact location of partition function zeros, a new method for statistical mechanics'. *Journal of Physics A: Mathematical and General*, 20(2):511, 1987.

[21] B. Haeupler and D. G. Harris. Parallel algorithms and concentration bounds for the Lovász local lemma via witness-dags. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1170–1187. SIAM, 2017.

[22] D. G. Harris. Lopsidependency in the moser-tardos framework: beyond the lopsided Lovász local lemma. *ACM Transactions on Algorithms (TALG)*, 13(1):17, 2016.

[23] D. G. Harris. Deterministic parallel algorithms for fooling polylogarithmic juntas and the Lovász local lemma. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1188–1203. SIAM, 2017.

[24] D. G. Harris and A. Srinivasan. A constructive algorithm for the Lovász local lemma on permutations. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 907–925. Society for Industrial and Applied Mathematics, 2014.

[25] N. J. Harvey and J. Vondrák. An algorithmic proof of the Lovász local lemma via resampling oracles. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 1327–1346. IEEE, 2015.

[26] K. Kolipaka, M. Szegedy, and Y. Xu. A sharper local lemma with improved applications. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 603–614. Springer, 2012.

[27] K. B. R. Kolipaka and M. Szegedy. Moser and Tardos meet Lovász. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 235–244. ACM, 2011.

[28] V. Kolmogorov. Commutativity in the algorithmic Lovász local lemma. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 780–787. IEEE, 2016.

[29] C. McDiarmid. Hypergraph colouring and the Lovász local lemma. *Discrete Mathematics*, 167:481–486, 1997.

[30] A. Moitra. Approximate counting, the Lovász local lemma and inference in graphical models. *arXiv preprint arXiv:1610.04317*, 2016.

[31] A. Moitra. Approximate counting, the lovasz local lemma, and inference in graphical models. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 356–369. ACM, 2017.

[32] M. Molloy and B. Reed. Further algorithmic aspects of the local lemma. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 524–529. ACM, 1998.

[33] R. A. Moser and G. Tardos. A constructive proof of the general Lovász local lemma. *Journal of the ACM (JACM)*, 57(2):11, 2010.

[34] W. Pegden. Highly nonrepetitive sequences: winning strategies from the local lemma. *Random Structures & Algorithms*, 38(1-2):140–161, 2011.

[35] W. Pegden. The lefthanded local lemma characterizes chordal dependency graphs. *Random Structures & Algorithms*, 41(4):546–556, 2012.

[36] W. Pegden. An extension of the moser–tardos algorithmic local lemma. *SIAM Journal on Discrete Mathematics*, 28(2):911–917, 2014.

[37] J. Radhakrishnan and A. Srinivasan. Improved bounds and algorithms for hypergraph two-coloring. In *Foundations of Computer Science, 1998. Proceedings. 39th Annual Symposium on*, pages 684–693. IEEE, 1998.

[38] M. R. Salavatipoury. A $(1+\epsilon)$-approximation algorithm for partitioning hypergraphs using a new algorithmic version of the Lovász local lemma.

[39] O. Sattath and I. Arad. A constructive quantum Lovász local lemma for commuting projectors. *Quantum Information & Computation*, 15(11-12):987–996, 2015.

[40] A. D. Scott and A. D. Sokal. The repulsive lattice gas, the independent-set polynomial, and the Lovász local lemma. *Journal of Statistical Physics*, 118(5):1151–1261, Mar 2005.

[41] A. D. Scott and A. D. Sokal. On dependency graphs and the lattice gas. *Combinatorics, Probability & Computing*, 15(1-2):253–279, 2006.

[42] J. B. Shearer. On a problem of Spencer. *Combinatorica*, 5(3):241–245, 1985.

[43] J. Spencer. Asymptotic lower bounds for Ramsey functions. *Discrete Mathematics*, 20:69–76, 1977.

[44] S. Todo. Transfer-matrix study of negative-fugacity singularity of hard-core lattice gas. *International Journal of Modern Physics C*, 10(04):517–529, 1999.

[45] D. Wood. The exact location of partition function zeros, a new method for statistical mechanics. *Journal of Physics A: Mathematical and General*, 18(15):L917, 1985.