

Article

A Meeting Point of Probability, Graphs, and Algorithms: The Lovász Local Lemma and Related Results—A Survey

András Faragó 

Department of Computer Science, The University of Texas at Dallas, 800 W. Campbell Rd., Richardson, TX 75080, USA; farago@utdallas.edu

Abstract: A classic and fundamental result, known as the Lovász Local Lemma, is a gem in the probabilistic method of combinatorics. At a high level, its core message can be described by the claim that weakly dependent events behave similarly to independent ones. A fascinating feature of this result is that even though it is a purely probabilistic statement, it provides a valuable and versatile tool for proving completely deterministic theorems. The Lovász Local Lemma has found many applications; despite being originally published in 1973, it still attracts active novel research. In this survey paper, we review various forms of the Lemma, as well as some related results and applications.

Keywords: Lovász Local Lemma; probabilistic method in combinatorics; probabilistic polynomial time algorithm



Citation: Faragó, A. A Meeting Point of Probability, Graphs, and Algorithms: The Lovász Local Lemma and Related Results—A Survey. *Algorithms* **2021**, *14*, 355. <https://doi.org/10.3390/a14120355>

Academic Editor: Jesper Jansson

Received: 14 November 2021

Accepted: 5 December 2021

Published: 8 December 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

“One of the most remarkable developments in Computer Science over the past 50 years has been the realization that allowing computers to toss coins can lead to algorithms that are more efficient, conceptually simpler and more elegant than their best known deterministic counterparts.”

Alistair Sinclair, University of California, Berkeley

1. Introduction

The *probabilistic method* in combinatorics is a very useful and powerful family of proof techniques; for a standard reference book, see Alon and Spencer [1]. As an important tool in this area, the result known as *Lovász Local Lemma (LLL)*, which was first published in 1973 by Erdős and Lovász [2], plays a fundamental role and is often referred to as a “gem” in this field.

A key (and quite fascinating) feature of the probabilistic method is that using probabilistic arguments, it allows proving completely deterministic claims. This is also the case for the Lovász Local Lemma. Specifically, it allows us to prove that a certain structure or property exists with positive probability in a random setting, which implies that the structure must indeed exist or the property must hold *deterministically*. In a sense, randomness and probability play here a catalyst role: they make possible the progress toward the goal, but eventually disappear in the end result. In a number of cases, however, no other way is known to achieve the same result, which is quite surprising.

In the next section, we describe a specific motivating problem and solve it with the Lovász Local Lemma, illustrating its strength.

2. A Motivating Problem and the Original LLL

Hypergraph 2-coloring. Let H be a hypergraph, in which every hyperedge has at least $k \geq 3$ vertices, and each hyperedge intersects at most d others. We ask that for what

values of k and d can it be guaranteed that such a hypergraph is 2-colorable—that is, we can color each vertex with one of two colors, such as red and blue, such that no hyperedge becomes monochromatic.

To describe a solution attempt, let E_1, \dots, E_n be the hyperedges in H . Consider a *random* coloring—that is, for each vertex, flip a coin independently and, depending on the outcome, color the vertex red or blue. Let A_i denote the event that E_i is monochromatic. What we aim to achieve is that no hyperedge is monochromatic, which is expressed by the event $C = \bar{A}_1 \dots \bar{A}_n$. If we can show $\Pr(C) > 0$, a 2-coloring of H must exist.

To find the probability of C , let us first estimate the probability of A_i . If E_i has ℓ_i vertices, then out of its 2^{ℓ_i} possible colorings, only two are monochromatic (the all-red and all-blue colorings). Therefore,

$$\Pr(A_i) = \Pr(E_i \text{ is monochromatic}) = 2/2^{\ell_i} \leq 2/2^k = 2^{1-k}, \tag{1}$$

where the inequality follows from $\ell_i \geq k$. Then, (1) implies $\Pr(\bar{A}_i) \geq 1 - 2^{1-k}$. Now, if the events A_i were independent, we could argue that

$$\Pr(C) = \Pr(\bar{A}_1 \dots \bar{A}_n) = \Pr(\bar{A}_1) \dots \Pr(\bar{A}_n) \geq (1 - 2^{1-k})^n > 0 \tag{2}$$

holds. It would mean that with positive probability no edge is monochromatic, so a 2-coloring of H must indeed exist. The problem with this argument is that the events A_1, \dots, A_n are *not* independent whenever the hyperedges overlap. Consequently, the equality $\Pr(\bar{A}_1 \dots \bar{A}_n) = \Pr(\bar{A}_1) \dots \Pr(\bar{A}_n)$ that we used in (2) generally does not hold.

This is the critical point where the LLL provides invaluable help. Before going into formal details, let us informally display the core message of the lemma:

Key message of LLL. *For any system A_1, \dots, A_n of events, in order for the conclusion $\Pr(\bar{A}_1 \dots \bar{A}_n) > 0$ to hold, it is enough if the events are “almost” independent, in the sense that each one of them depends only on a limited number of others, provided that $\Pr(A_i)$ is small enough.*

Specifically, the historically first version of the lemma, proved in the paper [2] published by Erdős and Lovász in 1973, is the following claim:

Lemma 1 (LLL—symmetric version, original form). *Let A_1, \dots, A_n be events such that each event is mutually independent of all the others, except at most $d \geq 1$ of them. If*

$$\Pr(A_i) \leq \frac{1}{4d} \quad (\forall i) \tag{3}$$

holds, then $\Pr(\bar{A}_1 \dots \bar{A}_n) > 0$.

Let us apply it to the above-described problem of hypergraph 2-coloring. Each event A_i in the example depends only on those events A_j for which $E_i \cap E_j \neq \emptyset$, since we pick the color of each vertex independently. As each hyperedge intersects at most d others, A_i may depend on at most d other events. We know from (1) that $\Pr(A_i) \leq 2^{1-k}$. Therefore, if

$$2^{1-k} \leq \frac{1}{4d} \tag{4}$$

holds, then $\Pr(A_i) \leq 2^{1-k} \leq \frac{1}{4d}$ follows, so the inequality (3) is satisfied. As a result, Lemma 1 implies $\Pr(\bar{A}_1 \dots \bar{A}_n) > 0$, which means that the desired 2-coloring must exist. Rearranging (4), we obtain the condition $4d \leq 2^{k-1}$, or equivalently, $d \leq 2^{k-3}$. Thus, by means of the LLL, we have proved the following:

Theorem 1. *Let H be a hypergraph in which each hyperedge has at least $k \geq 3$ vertices, and any hyperedge intersects at most $d \leq 2^{k-3}$ implies that H is 2-colorable.*

For example, if each hyperedge of H has at least 11 vertices, and intersects with at most 256 other hyperedges, then H is always 2-colorable, regardless of its structure. The reason is that with $d = 256$ and $k = 11$ the inequality $d \leq 2^{k-3}$ becomes $256 \leq 2^8$, which indeed holds.

Remark 1. *The factor 4 in the condition $4d \leq 2^{k-1}$ can be reduced to $e = 2.718\dots$ (the base of the natural logarithm), using stronger versions of the LLL (see Section 4), yielding the somewhat less demanding condition $d \leq 2^{k-1}/e$, instead of $d \leq 2^{k-3}$. This, however, does not change the principle of the proof.*

It is fascinating to note that applying the probabilistic claim of the LLL, we have obtained the completely deterministic result of Theorem 1. Furthermore, to the author's best knowledge, no proof is known of Theorem 1 without the use of the Lovász Local Lemma.

3. Naming Conventions

Before looking at stronger versions of the LLL, as well as related results, let us summarize some naming conventions that have become commonplace over the years.

- **Why “local?”** The “local” adjective in the name refers to the situation that each event is typically dependent only on a small number of others. This can be visualized by a *dependency graph*, in which each event is represented by a vertex, and the event is mutually independent of all events to which its vertex is not connected. In this graph representation, each event is dependent only on its neighborhood—that is, subject to local dependencies only. This explains the “local” in the name.

Remark 2. *The dependency graph is not the same as the graph that we could obtain by connecting two nodes whenever they are dependent. For example, if all the events are pairwise independent, but not mutually independent, then the latter graph would have no edge, wrongly suggesting that all the events are mutually independent.*

- **Naming the lemma after Lovász.** The first version of the LLL was published in the paper [2], jointly authored by Paul Erdős and László Lovász. Then, why is the lemma not called “Erdős–Lovász Local Lemma?” The reason is that Erdős insisted in every lecture he gave about the subject that this result was created by Lovász alone, even though they applied it together in their joint paper. This is mentioned in an interview with Lovász, see [3] (in Hungarian).
- **Symmetric vs. asymmetric versions.** The original version of the LLL (Lemma 1) is *symmetric* in the sense that each event is treated equally: they all have to satisfy the same probability bound (3). In the *asymmetric* variant (see Section 4.2), the events may satisfy different probability bounds.

4. Stronger Versions of LLL

4.1. Strengthening the Symmetric Version

A natural question in connection with the original symmetric LLL (Lemma 1) is this: how large can $\Pr(A_i)$ be, such that the conclusion $\Pr(\overline{A}_1 \dots \overline{A}_n) > 0$ is still guaranteed to hold? Note that a larger value of $\Pr(A_i)$ makes $\Pr(\overline{A}_i) = 1 - \Pr(A_i)$ smaller, so the nonemptiness of $\overline{A}_1 \dots \overline{A}_n$ becomes less likely. To formalize it, let us introduce a parameter $p > 0$ and replace the condition (3) with

$$\Pr(A_i) \leq p \quad (\forall i). \quad (5)$$

Then, the largest value of p for which the conclusion $\Pr(\overline{A}_1 \dots \overline{A}_n) > 0$ still remains valid provides us with the strongest form of the symmetric LLL. Note that the value of d is kept fixed.

Shearer [4] investigated the above question and computed a specific function $f(d)$, which is the supremum of p values in (5) for which the conclusion $\Pr(\bar{A}_1 \dots \bar{A}_n) > 0$ remains valid, for a fixed d :

Lemma 2 (Shearer’s lemma). For any fixed integer $d \geq 1$, let $f(d)$ be the supremum of p values, for which the the conclusion $\Pr(\bar{A}_1 \dots \bar{A}_n) > 0$ remains valid if p is used in place of $1/(4d)$ in (3). Then,

$$f(d) = \begin{cases} \frac{1}{2} & \text{if } d = 1 \\ \frac{(d-1)^{d-1}}{d^d} & \text{if } d \geq 2 \end{cases} \tag{6}$$

Since $f(d)$ is the supremum (not necessarily the maximum) of p values in (5) for which the conclusion $\Pr(\bar{A}_1 \dots \bar{A}_n) > 0$ holds, therefore, we can formulate the strongest version of the symmetric LLL this way:

Lemma 3 (Strongest symmetric LLL). Let A_1, \dots, A_n be events, such that each event is mutually independent of all the others, except at most $d \geq 1$ of them. Let $f(d)$ be given by (6). Then, if

$$\Pr(A_i) < f(d) \quad (\forall i)$$

holds, then $\Pr(\bar{A}_1 \dots \bar{A}_n) > 0$.

For example, if $d = 3$, then $1/(4d) \approx 0.0833$, but $f(3) = 4/27 \approx 0.148$, so in this case the latter allows an about 77.7% larger value for $\Pr(A_i)$, than (3). Since, however, the formula (6) for $f(d)$ is relatively complicated, one may look for a simpler bound. Spencer [5] proved that $p = \frac{1}{e^{(d+1)}}$ suffices in place of $\frac{1}{4d}$ in (3), where $e = 2.718\dots$ is the base of the natural logarithm. Harvey and Vondrák [6] further improved it to $\frac{1}{ed}$.

4.2. Asymmetric LLL

This is the more general version of the LLL, published by Spencer [5], in which it is allowed that different events can have different probability bounds. We again consider arbitrary events A_1, \dots, A_n , and express their dependencies by a dependency graph $G = (V, E)$. In this graph, $V = \{1, \dots, n\}$, and each event A_i is assumed mutually independent of the set of events $\{A_j \mid (i, j) \notin E\}$. In other words, each event is mutually independent of all its non-neighbors in the dependency graph.

Lemma 4 (Asymmetric LLL). Let A_1, \dots, A_n be a system of events with dependency graph $G = (V, E)$. Suppose there are real numbers $x_1, \dots, x_n \in [0, 1)$, such that

$$\Pr(A_i) \leq x_i \prod_{(i,j) \in E} (1 - x_j) \quad (\forall i).$$

Then,

$$\Pr(\bar{A}_1 \dots \bar{A}_n) \geq \prod_{i=1}^n (1 - x_i).$$

In particular, $\Pr(\bar{A}_1 \dots \bar{A}_n) > 0$ holds.

5. Further Application Examples

We have already presented an application (hypergraph 2-coloring) in Section 2. In this section, we present further interesting applications.

5.1. Disjoint Paths

Let G be a graph with two distinguished vertices $u \neq v$. Further, let $\mathcal{Q}_1, \dots, \mathcal{Q}_n$ be given sets of $u - v$ paths, such that each set contains m different paths. We would like to select a path $P_i \in \mathcal{Q}_i$ for each i , such that the selected paths are edge-disjoint. Under what conditions can this be done?

Assume the family of path systems is *diffuse* in the following sense: any $P \in \mathcal{Q}_i$ is edge-disjoint from all but at most k paths in \mathcal{Q}_j for every $j \neq i$. Let us now pick a path $P_i \in \mathcal{Q}_i$ uniformly at random from each \mathcal{Q}_i . Let E_{ij} denote the event that P_i and P_j are not edge-disjoint. Then, we have

$$\Pr(E_{ij}) \leq \frac{k}{m}$$

since, for any $P_i \in \mathcal{Q}_i$, the other path $P_j \in \mathcal{Q}_j$ can be chosen m different ways, and among these at most k can share an edge with P_i , by assumption. Further, any event $E_{i,j}$ is mutually independent of all events $E_{s,t}$ for which $\{i, j\} \cap \{s, t\} = \emptyset$ holds. The reason is that $E_{i,j}$ involves paths selected from \mathcal{Q}_i and \mathcal{Q}_j , and these are selected independently from the paths picked from \mathcal{Q}_s and \mathcal{Q}_t whenever there is no common index, i.e., $\{i, j\} \cap \{s, t\} = \emptyset$. This implies that $d < 2n$, where d is the maximum degree in the dependency graph of the $E_{i,j}$ events. Therefore, if we satisfy

$$\frac{k}{m} \leq \frac{1}{8n},$$

then the upper bound

$$\Pr(E_{i,j}) \leq \frac{1}{4d} \quad (\forall i, j)$$

holds, due to

$$\Pr(E_{i,j}) \leq \frac{k}{m} \leq \frac{1}{8n} = \frac{1}{4(2n)} \leq \frac{1}{4d}.$$

It means, the $E_{i,j}$ events satisfy the conditions of the LLL (Lemma 1), yielding

$$\Pr\left(\bigcap_{i,j} \bar{E}_{i,j}\right) > 0.$$

Consequently, there must exist a path system $P_i \in \mathcal{Q}_i, i = 1, \dots, n$, such that all the paths are edge-disjoint. Thus, by means of the LLL, we have proved the following result:

Theorem 2. *Let G be a graph with two distinguished vertices $u \neq v$, and let $\mathcal{Q}_1, \dots, \mathcal{Q}_n$ be sets of $u - v$ paths in G , such that each set contains m different paths. Assume that any $P \in \mathcal{Q}_i$ is edge-disjoint from all but at most k paths in \mathcal{Q}_j , for every $j \neq i$. Then, whenever $k \leq m/(8n)$ holds, it is possible to select a path from each \mathcal{Q}_i such that all the selected paths are edge-disjoint.*

5.2. k -SAT

Let us recall some well-known concepts about Boolean formulas. Such a formula is a CNF (Conjunctive Normal Form) formula if it is the conjunction (logical AND) of *clauses*, where each clause is the disjunction (logical OR) of *literals*. A literal is either a Boolean variable or its negation. A formula is a k -CNF formula if every clause contains k literals. We assume that the same variable cannot occur multiple times in a clause. The problem called k -SAT is this: given a k -CNF formula, is it satisfiable? This is a well-known NP-complete problem, but in some cases, the LLL allows us to quickly show that certain k -CNF formulas are satisfiable.

Let us say that two clauses in a k -CNF formula overlap, if there is a variable occurring in both (regardless of whether the variable is negated or not in the clauses). We can show via the LLL that if any clause overlaps with at most 2^{k-2} other clauses in a k -CNF formula Φ , then Φ is satisfiable.

Assign random truth values to the variables independently. Let $C_i, i = 1, \dots, n$, be the clauses and let A_i denote the event that C_i is not satisfied by the random truth assignment. There are 2^k possible truth assignments to the k literals in C_i , and only one of them makes

C_i false – the one in which all literals of C_i are false. Therefore, $\Pr(A_i) = 2^{-k}$. Furthermore, observe that each event A_i is mutually independent of the set of those A_j 's that correspond to clauses which do not overlap with C_i . Consequently, the maximum degree d in the dependency graph of the events A_1, \dots, A_n satisfies $d \leq 2^{k-2}$. Therefore, we can write

$$4 \Pr(A_i) d \leq 4 \cdot 2^{-k} 2^{k-2} = 1,$$

implying

$$\Pr(A_i) \leq \frac{1}{4d} \quad (\forall i).$$

Then, we can conclude from the LLL that $\Pr(\bar{A}_1 \dots \bar{A}_n) > 0$ holds, which means that at least one of the random truth assignments satisfies *all* clauses—that is, Φ is satisfiable. It is interesting to note that this does not depend on n , the number of clauses, and it is also independent of the number of variables. Thus, by means of the LLL, we have proved the following result:

Theorem 3. *If in a k -CNF formula Φ any clause overlaps with at most 2^{k-2} other clauses, then Φ is satisfiable.*

For example, if each clause contains $k = 12$ literals, and each clause overlaps with at most $2^{k-2} = 1024$ other clauses, then the formula is satisfiable, no matter how many clauses and variables it has.

5.3. Independent Sets in Graphs

Consider a graph with vertex set V and maximum degree Δ . Assume V is partitioned as $V = V_1 \cup \dots \cup V_r$, where each V_i satisfies $|V_i| \geq \ell\Delta$, for some positive integer ℓ . We ask: what value of ℓ guarantees that one can always select a vertex from each V_i , such that the selected vertices form an independent set, i.e., a set in which no two vertices are connected.

First, observe that we can assume $|V_i| = \ell\Delta$, since otherwise we can just remove vertices from V_i to achieve the equality. The reason is that if the claim holds with fewer vertices in V_i , then it certainly must hold in the original graph. Set $k = \ell\Delta$; this will be the size of each V_i , after possibly removing some vertices.

Now, pick a random vertex from each V_i . For each edge e , let A_e denote the event that e connects two of the randomly selected vertices. This can only happen if e connects vertices in two different sets V_i, V_j . Since $|V_i| = |V_j| = k$, $\Pr(A_e) = 1/k^2$ holds for any e due to the random vertex selection.

Observe now that if e has its endpoints in the sets A_i, A_j , then for another edge f , the event A_f can only depend on A_e if f has an endpoint in $A_i \cup A_j$. How many such edges f can exist? The size of $A_i \cup A_j$ is $2k$ and each vertex is adjacent to at most Δ edges. Therefore, the maximum degree in the dependency graph of the A_e events is at most $d = 2k\Delta$.

Let us now apply the LLL, in its simplest form, presented in Lemma 1. We need to satisfy the condition $\Pr(A_e) \leq 1/(4d)$. Due to $\Pr(A_e) = 1/k^2$, it is satisfied if

$$\frac{1}{k^2} \leq \frac{1}{4d} = \frac{1}{8k\Delta}$$

holds. From this, and from $k = \ell\Delta$, we obtain

$$\frac{1}{\ell^2 \Delta^2} \leq \frac{1}{8\ell \Delta^2}.$$

After simplification, this yields $\ell \geq 8$. Consequently, the Lovász Local Lemma implies that if $\ell \geq 8$, then

$$\Pr(\bar{A}_{e_1} \dots \bar{A}_{e_m}) > 0$$

holds, where e_1, \dots, e_m are the edges of the graph. This means, there is a positive probability that no edge connects two of the randomly selected vertices, so the special selection we

are looking for must indeed exist. Thus, by means of the LLL, we have proved the following result:

Theorem 4. *Let G be a graph with vertex set V and maximum degree Δ . Assume V is partitioned as $V = V_1 \cup \dots \cup V_r$, where each V_i satisfies $|V_i| \geq 8\Delta$. Then, it is always possible to select a vertex from each V_i such that the selected vertices form an independent set.*

Remark 3. The constant 8 can be reduced to $2e \approx 5.43656$ via the stronger versions of the symmetric LLL (see Section 4.1), but this does not change the principle of the proof.

5.4. Graph Coloring with Additional Constraints

The well-known task of graph coloring means that we need to assign a color to each vertex of the graph such that neighboring vertices have different colors. If a coloring satisfies this requirement, it is called a *proper coloring*.

The standard question regarding graph coloring is: how many colors are needed for a proper coloring? A simple basic fact is that in a graph with maximum degree Δ , it is always enough to use at most $\Delta + 1$ colors: index the colors by $1, 2, \dots$ and color the vertices one by one, always using the lowest indexed color that has not occurred among the already colored neighbors of the considered vertex. It is easy to see that this leads to the use of at most $\Delta + 1$ colors.

What if, however, we want a coloring that also satisfies some extra requirements? Such an extra requirement is that the frequency of color occurrences in the neighborhood of any vertex is limited, as defined below.

Definition 1. *Let $\beta \geq 1$ be an integer. A proper coloring is called β -frugal if no color appears more than β times in the neighborhood of any vertex.*

Hind, Molloy, and Reed [7] analyzed the number of colors needed for β -frugal colorings. They found a relationship between the maximum degree and the number of colors that suffice, but this is much more complex than the simple $\Delta + 1$ bound mentioned above for conventional coloring. Using the asymmetric Lovász Local Lemma, they proved the theorem below (we present a simplified version given by Sinclair [8]):

Theorem 5. *Let G be a graph. If the maximum degree Δ of G satisfies $\Delta \geq \beta^\beta$, then G has a β -frugal proper coloring using at most $16\Delta^{1+1/\beta}$ colors.*

5.5. Packet Scheduling in Networks

Consider a communication network, which is modeled by a directed graph. We have a given system P_1, \dots, P_r of directed paths in the network, and we want to send packets along these paths. The paths are edge-simple, which means that no path repeats any edge. Let us consider a synchronous communication model, in which for each time unit, one packet can traverse an edge. Assume a set of packets is given for each path. We would like to achieve that each packet is delivered on its respective path in the shortest possible time. Since the paths and packet sets are given, our only degree of freedom is to decide, whenever several packets compete for an edge, which one goes first, i.e., in which order they will traverse the edge (recall that only one can do it in one time unit). In other words, we are looking for the optimal packet scheduling policy. Furthermore, we would like to achieve it with a constant buffer size at every node.

Let us define some key parameters for this problem: *congestion* c is the maximum number of paths that share a directed edge; *dilation* d is the maximum path length that occurs in the path system. Let T be the smallest time in which each packet can be delivered, no matter how many packets wait for each path. The time for each individual packet is measured from the instant when it starts traversing the first edge on its path. A *lower bound* on T in the worst case is

$$T \geq \max\{c, d\}.$$

The justification for this lower bound is the following:

- A path can be d edge long, and traversing such a path clearly takes d time units, implying $T \geq d$.
- An edge can be included in c paths. If a packet arrives at the edge on each of these c paths at the same time, then one of them will suffer $c - 1$ time units delay, plus its own traversal time, resulting in $T \geq c$ in the worst case.
- Then, $T \geq d$ and $T \geq c$ together yield the worst case lower bound $T \geq \max\{c, d\}$.

A remarkable thing is that this essentially trivial lower bound can be achieved, up to a constant factor, no matter how many packets wait for each route. Furthermore, a constant size buffer suffices at every node. Specifically, by a sophisticated use of the Lovász Local Lemma, Leighton, Maggs, and Richa [9] proved the following theorem:

Theorem 6. *In the above model, there is always a packet schedule that achieves $O(c + d)$ delivery time for each packet, with a constant queue length at every node.*

Observe that $c + d \leq 2 \max\{c, d\}$; so $O(c + d)$ is indeed optimal up to a constant factor.

6. The Algorithmic Lovász Local Lemma

The original LLL (Lemma 1) and its variants presented in Section 4 are all *existence theorems*: they prove that there exists a realization that satisfies each of $\bar{A}_1 \dots \bar{A}_n$, but do not tell anything about how to find such a realization algorithmically.

In many cases, it is very natural to look for an algorithmic solution. For example, in the case presented in Example 3 (see Section 5.2), we have shown that in a k -CNF formula Φ , if every clause overlaps with at most 2^{k-2} other clauses, then Φ is satisfiable. In most applications, however, it is not enough just to know that a formula is satisfiable, we also want to find an actual satisfying truth assignment. The original LLL, as applied to this problem, does not provide it.

A naive approach would be to try rejection sampling to solve the above problem: repeatedly pick random truth assignments until one of them eventually satisfies all clauses. Since the LLL guarantees $\Pr(\bar{A}_1 \dots \bar{A}_n) > 0$, we must find a satisfying truth assignment after a finite number of trials, with probability 1. Unfortunately, this approach can be very inefficient, because the probability $\Pr(\bar{A}_1 \dots \bar{A}_n)$, although guaranteed to be positive by the LLL, is typically exponentially small. This would therefore lead to an exponential time algorithm, essentially not better than exhaustive search.

The first proof of the possibility of an efficient algorithmic LLL was published by Beck [10] in 1991. His algorithm, however, was somewhat technical, and required a stronger condition on neighborhoods than the original LLL. After a sequence of improvements, the most elegant solution was published by Moser and Tardos [11]. It is worth mentioning that the authors received the prestigious *Gödel Prize* in 2020 for this work.

The Moser–Tardos algorithm can be best presented via a formulation called the *variable version* of the LLL. Let X_1, \dots, X_m be mutually independent random variables. Denote the “bad” events that we want to avoid by A_1, \dots, A_n and let $p_i = \Pr(A_i)$ be the probability of A_i , $i = 1, \dots, n$. Let $\text{vbl}(A_i) \subseteq \{X_1, \dots, X_m\}$ denote the subset of variables on which A_i depends. In particular, if $\text{vbl}(A_i) \cap \text{vbl}(A_j) = \emptyset$, then A_i and A_j are independent. The dependency graph $G = (V, E)$ is defined by $V = \{1, \dots, n\}$ and $E = \{(i, j) \mid \text{vbl}(A_i) \cap \text{vbl}(A_j) \neq \emptyset\}$.

With the above framework, the Moser–Tardos algorithm is surprisingly simple and elegant, see Algorithm 1. We present the most concise version, found in Szegedy [12].

The correctness and complexity of the algorithm is captured by the following theorem, which implies that this is a probabilistic polynomial time algorithm.

Algorithm 1: RESAMPLE (Moser and Tardos).

- Assign random values independently to the variables X_1, \dots, X_m .
- While there is an $i, 1 \leq i \leq n$, such that \bar{A}_i is not satisfied by the current assignment, do:
 - Choose the smallest such i and **resample** all variables in $X_j \in \text{vbl}(A_i)$. That is, choose new random values independently for each $X_j \in \text{vbl}(A_i)$.
- Return the current variable assignment.

Theorem 7. *If the conditions of the asymmetric LLL hold, i.e., there are real numbers $x_1, \dots, x_n \in [0, 1)$, such that*

$$p_i \leq x_i \prod_{(i,j) \in E} (1 - x_j) \quad (\forall i),$$

then the algorithm resamples each $\text{vbl}(A_i)$ at most an expected number of $\frac{x_i}{1-x_i}$ times before finding an assignment that satisfies all the \bar{A}_i . Therefore, the expected overall number of resamplings is at most $\sum_{i=1}^n \frac{x_i}{1-x_i}$.

Furthermore, if the maximum degree of the dependency graph is D , and $p_i \leq \frac{1}{e^{(D+1)}} \quad (\forall i)$ holds, then $\text{vbl}(A_i)$ is resampled at most an expected number of $1/D$ times, and the total expected number of resamplings is bounded by n/D .

While the algorithm itself is surprisingly straightforward and elegant, the proof of the above Theorem is far from simple, see Moser and Tardos [11]. It is interesting to note that the RESAMPLE algorithm is similar to the WalkSAT heuristic, which is used to solve general Boolean satisfiability problems, see, e.g., Hoos and Stützle [13].

7. Outlook

Let us briefly list some other recent research directions that are being pursued in connection with the Lovász Local Lemma and its algorithmic version, without going into details.

- *Derandomization* of the Moser–Tardos algorithm is possible, leading to a deterministic polynomial time algorithm. This was already addressed by Moser and Tardos in [11], and further developed by Chandrasekaran, Goyal, and Haeupler [14].
- Some recent papers deal with *sampling and counting* problems related to the LLL, see, e.g., Jain, Pham, and Vuong [15].
- *Approximate counting* is also considered in the context, such as approximately counting the satisfying truth assignments of a CNF formula, see Moitra [16].
- A *quantum version* of the LLL has also been introduced, referred to as *Quantum Lovász Local Lemma*, see Ambainis, Kempe, and Sattath [17] and He, Li, Sun, and Zhang [18]. The key difference between the Quantum Lovász Local Lemma and the classical LLL is that in the quantum version, the events are substituted with subspaces, and the event probabilities are substituted with subspace dimensions. This makes it more suitable for quantum computing applications.
- The LLL is also being used to analyze *parallel and distributed algorithms*, see, e.g., Chang, He, Li, Pettie, and Uitto [19].
- There are attempts to extend the LLL to an infinite setting, see Bernshteyn [20].

8. Conclusions

We have surveyed a classic and fundamental result, known as the Lovász Local Lemma. We reviewed some variants, related results, and applications, as well as the algorithmic version. We were delighted with the intriguing feature that a purely probabilistic claim can be used to prove completely deterministic statements.

As a closing conclusion, we can claim that in a high-level, abstract sense, the Lovász Local Lemma also provides efficiency in a way that is somewhat reminiscent of the fact that general graph problems often become more tractable in bounded degree graphs. Of course, the LLL case is more complicated, since it involves not only neighborhood sizes, but also probabilities, but it still deals with bounded degree style structures. As an example, let us refer to the k -SAT problem (see Section 5.2) that becomes solvable in polynomial time, whenever the clause overlaps are limited.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Conflicts of Interest: The author declares no conflict of interest.

References

1. Alon, N.; Spencer, J.H. *The Probabilistic Method*, 4th ed.; Wiley: Hoboken, NJ, USA, 2016.
2. Erdős, P.; Lovász, L. Problems and Results on 3-chromatic Hypergraphs and Some Related Questions. In *Infinite and Finite Sets*; Hajnal, A., Rado, R., Sós, V.T., Eds.; North-Holland Publishing Co.: Amsterdam, The Netherlands, 1973; Volume 2, pp. 609–627.
3. Laczkovich, M. A Random Walk in and Around Mathematics—Interview with László Lovász. *Magyar Tudomány (Hung. Sci.)* **2021**, *182*, 1108–1123.
4. Shearer, J.B. On a Problem of Spencer. *Combinatorica* **1985**, *5*, 241–245. [[CrossRef](#)]
5. Spencer, J. Asymptotic Lower Bounds for Ramsey Functions. *Discret. Math.* **1997**, *20*, 69–76. [[CrossRef](#)]
6. Harvey, N.J.A.; Vondrák, J. Short Proofs for Generalizations of the Lovász Local Lemma: Shearer’s Condition and Cluster Expansion. *arXiv* **2017**, arXiv:1711.06797.
7. Hind, H.; Molloy, M.; Reed, B. Coloring a Graph Frugally. *Combinatorica* **1997**, *17*, 469–482. [[CrossRef](#)]
8. Sinclair, A. CS271 Randomness & Computation: Spring 2020, Lecture 22. Available online: <https://people.eecs.berkeley.edu/~sinclair/cs271/s20.html> (accessed on 1 December 2021).
9. Leighton, F.T.; Maggs, B.M.; Richa, A. Fast Algorithms for Finding $O(\text{Congestion} + \text{Dilation})$ Packet Routing Schedules. *Combinatorica* **1999**, *19*, 375–401. [[CrossRef](#)]
10. Beck, J. An Algorithmic Approach to the Lovász Local Lemma. *Random Struct. Algorithms* **1991**, *2*, 343–366. [[CrossRef](#)]
11. Moser, R.A.; Tardos, G. A Constructive Proof of the General Lovász Local Lemma. *J. ACM* **2010**, *57*, 1–15. [[CrossRef](#)]
12. Szegedy, M. The Lovász Local Lemma—A Survey. In Proceedings of the 8th International Computer Science Symposium in Russia (CSR 2013), Ekaterinburg, Russia, 25–29 June 2013; pp. 1–11.
13. Hoos H.; Stützle, T. Local Search Algorithms for SAT: An Empirical Evaluation. *J. Autom. Reason.* **2000**, *24*, 421–481. [[CrossRef](#)]
14. Chandrasekaran, K.; Goyal, N.; Haeupler, B.B. Deterministic Algorithms for the Lovász Local Lemma. *SIAM J. Comput.* **2013**, *42*, 2132–2155. [[CrossRef](#)]
15. Jain, V.; Pham, H.T.; Vuong, T.D. Towards the Sampling Lovász Local Lemma. *arXiv* **2020**, arXiv:2011.12196v1.
16. Moitra, A. Approximate Counting, the Lovász Local Lemma, and Inference in Graphical Models. *J. ACM* **2019**, *66*, 1–25. [[CrossRef](#)]
17. Ambainis, A.; Kempe, J.; Sattath, O. A Quantum Lovász Local Lemma. In Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC’10), Cambridge, MA, USA, 6–8 June 2010; pp. 151–160.
18. He, K.; Li, Q.; Sun, X.; Zhang, J. Quantum Lovász Local Lemma: Shearer’s Bound is Tight. In Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC’19), Phoenix, AZ, USA, 23–26 June 2019; pp. 461–472.
19. Chang, Y.-J.; He, Q.; Li, W.; Pettie, S.; Uitto, J. Distributed Edge Coloring and a Special Case of the Constructive Lovász Local Lemma. *ACM Trans. Algorithms* **2020**, *16*, 1–51. [[CrossRef](#)]
20. Bernshteyn, A. Measurable Versions of the Lovász Local Lemma and Measurable Graph Colorings. *Adv. Math.* **2019**, *353*, 153–223. [[CrossRef](#)]