The Lovász Local Lemma: constructive aspects, stronger variants and the hard core model

Jan Vondrák[1]

[1]Dept. of Mathematics
Stanford University

joint work with Nick Harvey (UBC)

The Lovász Local Lemma

Theorem (Symmetric LLL, Lovász $\sim$ 1975)

If $E_1, \ldots, E_n$ are events on a probability space Ω such that

- *Each event is independent of all but d other events*
- *The probability of each event is at most $\frac{1}{e(d+1)}$ ($e = 2.718..$)*

then

$$\Pr[\bigcap_{i=1}^{n} \overline{E_i}] > 0.$$

"Needle in a haystack" problem:

1. LLL implies that it is possible to avoid all events $E_1, \ldots, E_n$
2. but the probability of $\bigcap_{i=1}^{n} \overline{E_i}$ could be exponentially small

Example: the r-partite Turán problem

Consider an r-partite graph, at least $\rho|V_i||V_j|$ edges between every pair (V_i, V_j).

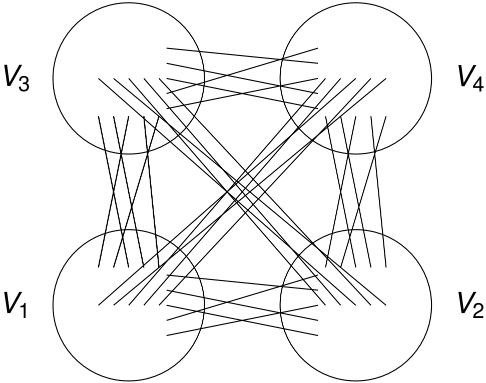

Question: at what density ρ must G contain K_r?

Application of the LLL

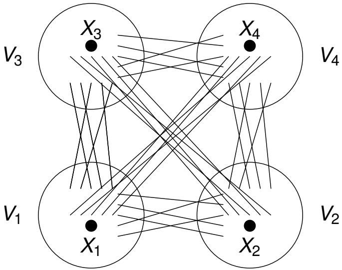

X_i = random vertex in V_i
E_{ij} = the event that $(X_i, X_j) \notin E$

We want: $(X_1, \ldots, X_r)$
such that no event E_{ij} occurs.

Application of the LLL

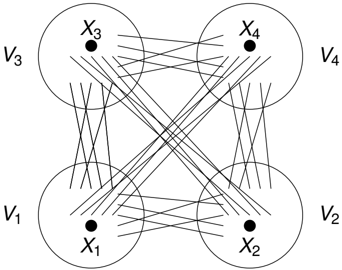

X_i = random vertex in V_i
E_{ij} = the event that $(X_i, X_j) \notin E$

We want: $(X_1, \ldots, X_r)$
such that no event E_{ij} occurs.

Parameters: $\Pr[E_{ij}] = 1 - \rho$, $d = 2(r-1)$
(dependencies only between $E_{ij}, E_{i'j}$ sharing an index).

Application of the LLL

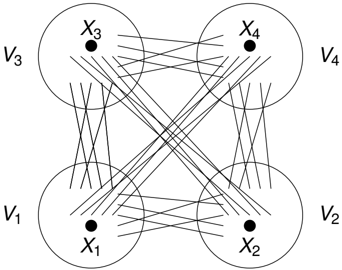

X_i = random vertex in V_i
E_{ij} = the event that $(X_i, X_j) \notin E$

We want: $(X_1, \ldots, X_r)$
such that no event E_{ij} occurs.

Parameters: $\Pr[E_{ij}] = 1 - \rho$, $d = 2(r-1)$
(dependencies only between $E_{ij}, E_{i'j}$ sharing an index).

LLL implies: If $\rho \geq 1 - \frac{1}{e(2r-1)}$ then G contains a K_r.

Application of the LLL

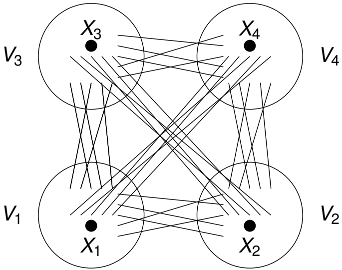

X_i = random vertex in V_i
E_{ij} = the event that $(X_i, X_j) \notin E$

We want: $(X_1, \ldots, X_r)$
such that no event E_{ij} occurs.

Parameters: $\Pr[E_{ij}] = 1 - \rho$, $d = 2(r-1)$
(dependencies only between $E_{ij}, E_{i'j}$ sharing an index).

LLL implies: If $\rho \geq 1 - \frac{1}{e(2r-1)}$ then G contains a K_r.

(Roughly correct: There is a graph with $\rho = 1 - \frac{1}{r-1}$ without a K_r.)

The General (asymmetric) Lovász Local Lemma

Theorem (General LLL)

If $E_1, \ldots, E_n$ are events with a "dependency graph", $\Gamma(i)$ = neighborhood of i, so that

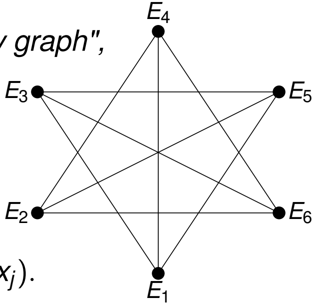

- *Each event E_i is independent of all the events $E_j, j \notin \Gamma(i) \cup \{i\}$*
- *There are $x_i \in (0, 1)$ such that*

$$\Pr[E_i] \leq x_i \prod_{j \in \Gamma(i)} (1 - x_j).$$

Then

$$\Pr[\bigcap_{i=1}^{n} \overline{E_i}] \geq \prod_{i=1}^{n} (1 - x_i).$$

(Symmetric variant can be obtained by setting $x_i = e \cdot Pr[E_i]$.)

Shearer's Lemma

(“optimal form of the local lemma")

For events $E_1, \ldots, E_n$ with probabilities $p_1, \ldots, p_n$ and a dependency graph G, define

$$q_S(p_1, \ldots, p_n) = \sum_{\text{indep. } I \subseteq S} (-1)^{|I|} \prod_{i \in I} p_i$$

(alternating-sign independence polynomial of the dependency graph).

Lemma (Shearer 1985)

If $\forall S \subseteq [n], q_S(p_1, \ldots, p_n) > 0$, *then*

$$\Pr[\bigcap_{i=1}^{n} \overline{E_i}] \geq q_{[n]}(p_1, \ldots, p_n).$$

(If not, then $\Pr[\bigcap_{i=1}^{n} \overline{E_i}]$ could be 0.)

Connection with statistical physics

[Scott-Sokal 2005]

Shearer's Lemma is closely related to the *hard core model of repulsive gas* in statistical physics.

Model:
particles on a graph G,
two particles never adjacent;
activity parameters w_i.

$\Pr[I] \sim \prod_{i \in I} w_i$ if I independent.

Partition function:
$Z(\mathbf{w}) = \sum_{\text{indep. } I \subseteq V} \prod_{i \in I} w_i$.

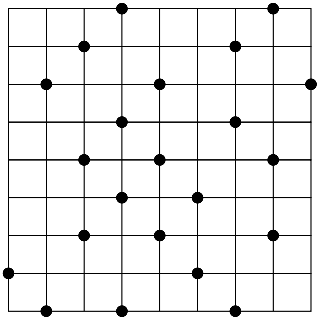

Fact: $\log Z(\mathbf{w})$ *has an alternating-sign Taylor series around* 0
("Mayer expansion").

Hard core model vs. Shearer's Lemma

[Scott-Sokal 2005] The following are equivalent:

1. Mayer expansion of $\log Z(\mathbf{w})$ is convergent for $|w_i| \leq R_i$.
2. $Z(-\lambda \mathbf{R}) > 0$ for all $0 \leq \lambda \leq 1$.
3. $Z_S(-\mathbf{R}) > 0$ for all subsets of vertices S, where

$$Z_S(\mathbf{w}) = \sum_{\text{indep. } I \subseteq S} \prod_{i \in I} w_i.$$

Hard core model vs. Shearer's Lemma

[Scott-Sokal 2005] The following are equivalent:

1. Mayer expansion of $\log Z(\mathbf{w})$ is convergent for $|w_i| \leq R_i$.
2. $Z(-\lambda \mathbf{R}) > 0$ for all $0 \leq \lambda \leq 1$.
3. $Z_S(-\mathbf{R}) > 0$ for all subsets of vertices S, where

$$Z_S(\mathbf{w}) = \sum_{\text{indep. } I \subseteq S} \prod_{i \in I} w_i.$$

Note: $Z_S(-\mathbf{p}) = q_S(\mathbf{p})$ are the quantities in Shearer's Lemma (whose positivity implies that all events can be avoided).

Hard core model vs. Lovász Local Lemma

Let $\Gamma(i)$ = neighborhood of i, and $\Gamma^+(i) = \{i\} \cup \Gamma(i)$.

Various sufficient conditions for the convergence of $\log Z(\mathbf{w})$ have been investigated.

- [Dobrushin 1996]
  If $w_i \leq y_i / \prod_{j \in \Gamma^+(i)} (1 + y_j)$ for some $y_i > 0$,
  then the Mayer expansion for $\log Z(\mathbf{w})$ converges.
  Corresponds exactly to the LLL (substitute $y_i = \frac{x_i}{1-x_i}$).

Hard core model vs. Lovász Local Lemma

Let $\Gamma(i)$ = neighborhood of i, and $\Gamma^+(i) = \{i\} \cup \Gamma(i)$.

Various sufficient conditions for the convergence of $\log Z(\mathbf{w})$ have been investigated.

- [Dobrushin 1996]
  If $w_i \leq y_i / \prod_{j \in \Gamma^+(i)} (1 + y_j)$ for some $y_i > 0$,
  then the Mayer expansion for $\log Z(\mathbf{w})$ converges.
  Corresponds exactly to the LLL (substitute $y_i = \frac{x_i}{1-x_i}$).
- [Fernandez-Procacci 2007]
  If $w_i \leq y_i / \sum_{\text{indep. } I \subseteq \Gamma^+(i)} \prod_{i \in I} y_i$ for some $y_i > 0$,
  then the Mayer expansion for $\log Z(\mathbf{w})$ converges.
  New criterion — previously unknown to combinatorialists.

The Cluster Expansion Lemma

Theorem (Bissacot-Fernandez-Procacci-Scoppola 2011)

If $E_1, \ldots, E_n$ are events with a dependency graph G,

- *Each event E_i is independent of its non-neighbor events.*
- *There are $y_i > 0$ such that*

$$\Pr[E_i] \le \frac{y_i}{\sum_{\text{indep. } I \subseteq \Gamma^+(i)} \prod_{i \in I} y_i}.$$

(To compare: in LLL, we sum up over ***all*** *subsets $I \subseteq \Gamma^+(i)$.)*

Then

$$\Pr[\bigcap_{i=1}^{n} \overline{E_i}] > 0.$$

(Analytic Proof.)

Combinatorial proof of Cluster Expansion

[Harvey-V. '15]

Define: $\overline{P}_S = \Pr[\bigcap_{i \in S} \overline{E_i}]$, $Y_S = \sum_{\text{indep. } I \subseteq S} \prod_{i \in I} y_i$.
We assume: $\Pr[E_i] \leq y_i / Y_{\Gamma^+(i)}$.

Combinatorial proof of Cluster Expansion

[Harvey-V. '15]

Define: $\overline{P}_S = \Pr[\bigcap_{i \in S} \overline{E_i}]$, $Y_S = \sum_{\text{indep. } I \subseteq S} \prod_{i \in I} y_i$.
We assume: $\Pr[E_i] \leq y_i / Y_{\Gamma^+(i)}$.

Recursive bounds:
$\overline{P}_S = \Pr[\bigcap_{i \in S-a} \overline{E_i}] - \Pr[E_a \wedge \bigcap_{i \in S-a} \overline{E_i}] \geq \overline{P}_{S-a} - p_a \overline{P}_{S \setminus \Gamma^+(a)}$,
$Y_{T+a} = Y_T + y_a Y_{T \setminus \Gamma^+(a)} \geq Y_T + p_a Y_{T \cup \Gamma^+(a)}$.

Combinatorial proof of Cluster Expansion

[Harvey-V. '15]

Define: $\overline{P}_S = \Pr[\bigcap_{i \in S} \overline{E_i}]$, $Y_S = \sum_{\text{indep. } I \subseteq S} \prod_{i \in I} y_i$.
We assume: $\Pr[E_i] \leq y_i / Y_{\Gamma^+(i)}$.

Recursive bounds:
$\overline{P}_S = \Pr[\bigcap_{i \in S-a} \overline{E_i}] - \Pr[E_a \wedge \bigcap_{i \in S-a} \overline{E_i}] \geq \overline{P}_{S-a} - p_a \overline{P}_{S \setminus \Gamma^+(a)}$,
$Y_{T+a} = Y_T + y_a Y_{T \setminus \Gamma^+(a)} \geq Y_T + p_a Y_{T \cup \Gamma^+(a)}$.

We claim, by induction,

$$\frac{\overline{P}_S}{\overline{P}_{S-a}} \geq \frac{Y_{\overline{S}}}{Y_{\overline{S-a}}} > 0.$$

Combinatorial proof of Cluster Expansion

[Harvey-V. '15]

Define: $\overline{P}_S = \Pr[\bigcap_{i \in S} \overline{E_i}]$, $Y_S = \sum_{\text{indep. } I \subseteq S} \prod_{i \in I} y_i$.
We assume: $\Pr[E_i] \leq y_i / Y_{\Gamma^+(i)}$.

Recursive bounds:
$\overline{P}_S = \Pr[\bigcap_{i \in S-a} \overline{E_i}] - \Pr[E_a \wedge \bigcap_{i \in S-a} \overline{E_i}] \geq \overline{P}_{S-a} - p_a \overline{P}_{S \setminus \Gamma^+(a)}$,
$Y_{T+a} = Y_T + y_a Y_{T \setminus \Gamma^+(a)} \geq Y_T + p_a Y_{T \cup \Gamma^+(a)}$.

We claim, by induction,

$$\frac{\overline{P}_S}{\overline{P}_{S-a}} \geq \frac{Y_{\overline{S}}}{Y_{\overline{S-a}}} > 0.$$

Proof:

$$\frac{\overline{P}_S}{\overline{P}_{S-a}} \geq 1 - p_a \frac{\overline{P}_{S \setminus \Gamma^+(a)}}{\overline{P}_{S-a}} \geq 1 - p_a \frac{Y_{\overline{S} \cup \Gamma^+(a)}}{Y_{\overline{S}+a}} \geq \frac{Y_{\overline{S}}}{Y_{\overline{S-a}}}.$$

Hierarchy of the Local Lemmas

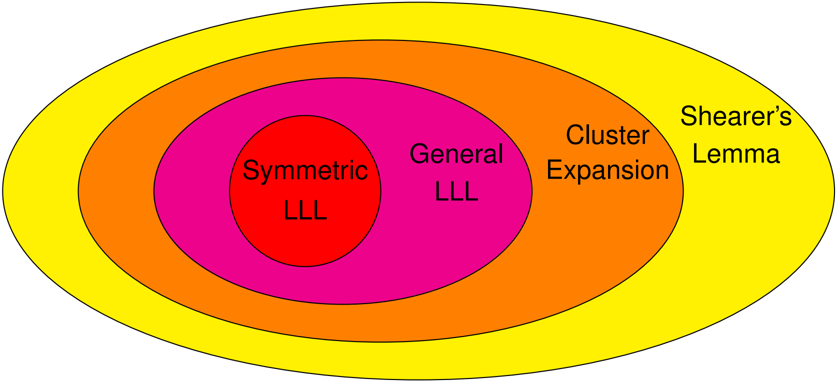

Application of Cluster Expansion to r-partite Turán

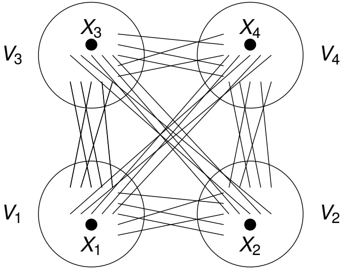

X_i = random vertex in V_i
E_{ij} = the event that $(X_i, X_j) \notin E$
$p_{ij} = \Pr[E_{ij}] = 1 - \rho$.

Dependency graph G = line graph of K_r. Neighborhood of E_{ij}: *two cliques*, events incident to i and events incident to j.

$$\sum_{\text{indep. } I \subseteq \Gamma^+(ij)} \prod_{(i'j') \in I} y_{i'j'} \leq (1 + \sum_{j'=1}^{r} y_{ij'})(1 + \sum_{i'=1}^{r} y_{i'j}) = (1 + ry)^2$$

(when all y_{ij} equal)

Set $y = \frac{1}{r}$: $\quad \frac{y}{(1+ry)^2} = \frac{1}{4r}$.

$\Rightarrow$ G always contains a K_r for $\rho \geq 1 - \frac{1}{4r}$. (improvement from $1 - \frac{1}{2er}$)

Application of Shearer's Lemma

[Csikváry-Nagy 2012]

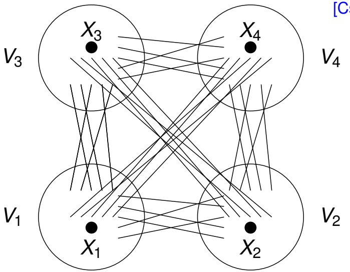

$$X_i = \text{random vertex in } V_i$$
$$E_{ij} = \text{the event that } (X_i, X_j) \notin E$$
$$p_{ij} = \Pr[E_{ij}] = 1 - \rho.$$

Dependency graph G = line graph of K_r
(events $E_{ij}, E_{i'j'}$ dependent if they share an index).

Independence polynomial $q(p)$ of G
= matching polynomial of K_r = *the Hermite polynomial*.

Application of Shearer's Lemma

[Csikváry-Nagy 2012]

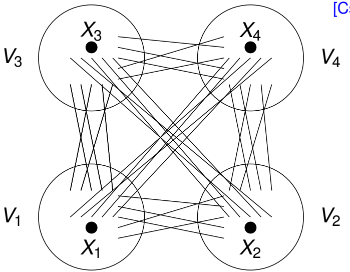

X_i = random vertex in V_i
E_{ij} = the event that $(X_i, X_j) \notin E$
$p_{ij} = \Pr[E_{ij}] = 1 - \rho$.

Dependency graph G = line graph of K_r
(events $E_{ij}, E_{i'j'}$ dependent if they share an index).

Independence polynomial $q(p)$ of G
= matching polynomial of K_r = *the Hermite polynomial*.

Roots of $q(p)$ well understood: minimum positive root $\geq \frac{1}{4(r-2)}$.
$\Rightarrow$ G always contains a K_r for $\rho \geq 1 - \frac{1}{4(r-2)}$.

Tight bound for the r-partite Turán problem?

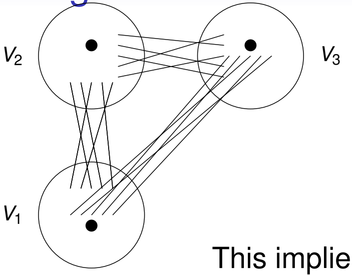

Shearer's Lemma:

For K_3, the matching polynomial is

$$q_3(p) = 1 - 3p.$$

Minimum root $p_0 = 1/3$.

This implies a K_3 subgraph for density $\rho \geq 2/3$.

Tight bound for the r-partite Turán problem?

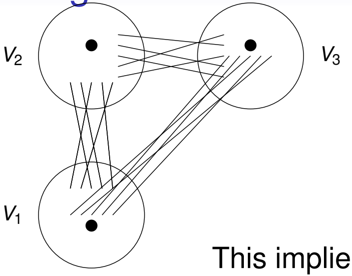

Shearer's Lemma:
For K_3, the matching polynomial is
$$q_3(p) = 1 - 3p.$$

Minimum root $p_0 = 1/3$.
This implies a K_3 subgraph for density $\rho \geq 2/3$.

But this is not tight: [Bondy-Shen-Thomassé-Thomassen 2006]
The optimal density for K_3 is $\rho^* = \frac{-1+\sqrt{5}}{2}$.

Open question: What is the optimal density that guarantees the appearance of K_r in an r-partite graph, for $r \geq 4$?
(roughly between $1 - \frac{1}{4r}$ and $1 - \frac{1}{2r}$)

The non-constructive aspect of the LLL

The proof of LLL is essentially non-constructive: $\Pr[\bigcap_{i=1}^{n} \overline{E_i}]$ is proved to be positive, but it could be exponentially small.

How do we find a state $\omega \in \bigcap_{i=1}^{n} \overline{E_i}$ efficiently, given an instance where the LLL applies?

The non-constructive aspect of the LLL

The proof of LLL is essentially non-constructive: $\Pr[\bigcap_{i=1}^{n} \overline{E_i}]$ is proved to be positive, but it could be exponentially small.

How do we find a state $\omega \in \bigcap_{i=1}^{n} \overline{E_i}$ efficiently, given an instance where the LLL applies?

Example:
Given an r-partite graph on n vertices, density of each pair $\rho \geq 1 - \frac{1}{4r}$. Can you find a K_r subgraph in $poly(n, r)$ time?

The Moser-Tardos framework

- Independent random variables $X_1, \ldots, X_m$.
- "Bad events" $E_1, \ldots, E_n$.
- Event E_i depends on variables $var(E_i)$.
- A dependency graph G:
  i—j iff $var(E_i) \cap var(E_j) \neq \emptyset$.
- There are $x_1, \ldots, x_n \in (0, 1)$ s.t.
  $\forall i;\ \ \Pr[E_i] \leq x_i \prod_{j \in \Gamma(i)}(1 - x_j)$.
  (asymmetric LLL condition)

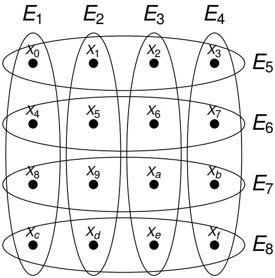

Moser-Tardos Algorithm:
Start with random variables $X_1, \ldots, X_m$. As long as some event E_i occurs, resample the variables in $var(E_i)$.

The Moser-Tardos framework

- Independent random variables $X_1, \dots, X_m$.
- "Bad events" $E_1, \dots, E_n$.
- Event E_i depends on variables $var(E_i)$.
- A dependency graph G:
  i—j iff $var(E_i) \cap var(E_j) \neq \emptyset$.
- There are $x_1, \dots, x_n \in (0,1)$ s.t.
  $\forall i;\ \ \Pr[E_i] \leq x_i \prod_{j \in \Gamma(i)} (1 - x_j)$.
  (asymmetric LLL condition)

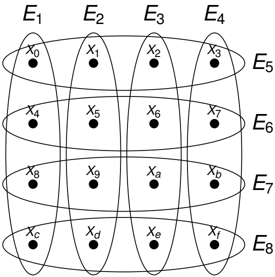

Moser-Tardos Algorithm:
Start with random variables $X_1, \dots, X_m$. As long as some event E_i occurs, resample the variables in $var(E_i)$.

Theorem (Moser-Tardos '08)

This algorithm finds $\omega \in \bigcap_{i=1}^{n} \overline{E_i}$ after $\sum_{i=1}^{n} \frac{x_i}{1-x_i}$ resampling operations (in expectation).

Beyond independent random variables

The LLL gives interesting applications also in spaces with more structure:

- permutations
- Hamiltonian cycles
- matchings
- trees

Beyond independent random variables

The LLL gives interesting applications also in spaces with more structure:

- permutations
- Hamiltonian cycles
- matchings
- trees

Follow-up work:

- [Kolipaka-Szegedy '11] extension of Moser-Tardos to Shearer's setting.
- [Harris-Srinivasan '14] handle applications with random permutations.
- [Achlioptas-Iliopoulos '14] general approach based on random walks; handle Hamiltonian cycles, matchings.

"Algorithmic proof" of the LLL?

We would like to have:

- given a probability space with events satisfying the LLL conditions, a (randomized) procedure that quickly finds $\omega \in \bigcap_{i=1}^{n} \overline{E_i}$.

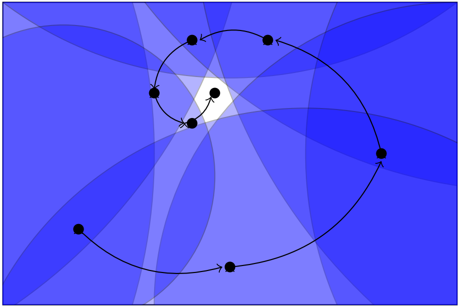

Ω

Our Main Result

"Algorithmic proof" of Shearer's Lemma:

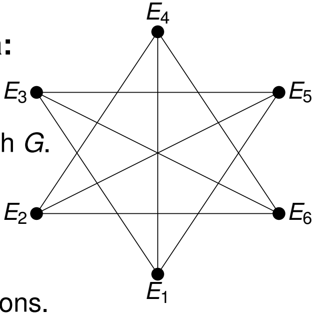

- Arbitrary probability space Ω.
- Events $E_1, \ldots, E_n$ with a dependency graph G.
- Each E_i independent of non-neighbors (or more generally, "positively associated" with non-neighbors)
- $p_i = (1 + \epsilon) \Pr[E_i]$ satisfy Shearer's conditions.

Our Main Result

"Algorithmic proof" of Shearer's Lemma:

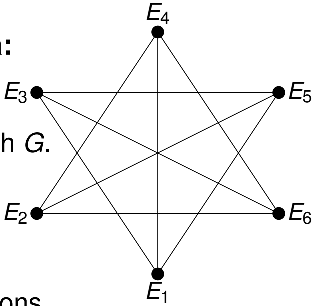

- Arbitrary probability space Ω.
- Events $E_1, \ldots, E_n$ with a dependency graph G.
- Each E_i independent of non-neighbors (or more generally, "positively associated" with non-neighbors)
- $p_i = (1 + \epsilon) \Pr[E_i]$ satisfy Shearer's conditions.

Theorem (Harvey-V. '15)

There is a randomized procedure which finds $\omega \in \bigcap_{i=1}^{n} \overline{E_i}$ *under these assumptions after* $O(\frac{n}{\epsilon} \log \frac{1}{\epsilon})$ *"resampling operations" w.h.p.*

Resampling operations

Assume a space Ω with a probability measure μ, and events $E_1, \ldots, E_n$ with a neighborhood structure denoted $\Gamma(i)$.

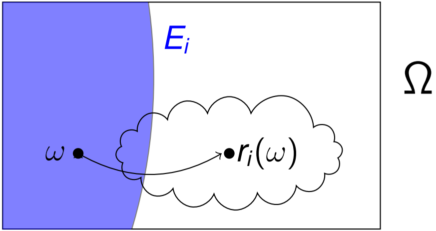

Definition: A **resampling operation** r_i for event E_i is a random $r_i(\omega) \in \Omega$ for each $\omega \in \Omega$, such that

1. If ω has distribution μ conditioned on $E_i \Rightarrow r_i(\omega)$ has distribution μ.
   (removes conditioning on E_i)
2. If $k \notin \Gamma^+(i)$ and $\omega \notin E_k \Rightarrow r_i(\omega) \notin E_k$.
   (does not cause non-neighbor events)

Why should resampling operations exist?

Lemma (Harvey-V. '15)

Resampling operations for events $E_1, \ldots, E_n$ w.r.t. G exist whenever each E_i is independent of its non-neighbor events.

Why should resampling operations exist?

Lemma (Harvey-V. '15)

Resampling operations for events $E_1, \ldots, E_n$ w.r.t. G exist whenever each E_i is independent of its non-neighbor events.

More generally: Resampling operations exist if and only if each E_i is "positively associated" with its non-neighbor events:

$$\mathbb{E}[Z \mid E_i] \geq \mathbb{E}[Z]$$

for every monotonic function Z of $(E_j : j \notin \Gamma^+(i))$.

Remark:

necessary to handle permutations and matchings; not for trees.

The algorithm

Our algorithm:
Sample ω from μ.
While any violated events exist, repeat:

- $J \leftarrow \emptyset$
- As long as $\exists j \notin \Gamma^+(J)$, E_j occurs

$$\{$$
$$\omega \leftarrow r_j(\omega), \quad \textbf{(resample } E_i\textbf{)}$$
$$J \leftarrow J \cup \{j\}$$
$$\}$$

Note: In each iteration we resample an independent set of events J.
In the next iteration, all violated events are in $\Gamma^+(J) = J \cup \Gamma(J)$.
($\Gamma(J)$ = neighbors of J in the dependency graph G)

Analysis of our algorithm

Def.: $Stab = \{(I_1, I_2, \ldots, I_t) : I_i \in \mathrm{Ind(G)} \setminus \{\emptyset\}, I_{i+1} \subseteq \Gamma^+(I_i)\}$.

Coupling lemma: The probability that the algorithm resamples a sequence of independent sets $(I_1, I_2, \ldots, I_t) \in Stab$ is at most

$$p(I_1, \ldots, I_t) = \prod_{s=1}^{t} \prod_{i \in I_s} p_i \qquad (p_i = \Pr_{\mu}[E_i]),$$

$$\mathbb{E}[\#\text{iterations}] \leq \sum_{t=0}^{\infty} \sum_{(I_1,\ldots,I_t) \in Stab} p(I_1, \ldots, I_t).$$

Analysis of our algorithm

Def.: $Stab = \{(I_1, I_2, \ldots, I_t) : I_i \in \mathrm{Ind}(G) \setminus \{\emptyset\}, I_{i+1} \subseteq \Gamma^+(I_i)\}$.

Coupling lemma: The probability that the algorithm resamples a sequence of independent sets $(I_1, I_2, \ldots, I_t) \in Stab$ is at most

$$p(I_1, \ldots, I_t) = \prod_{s=1}^{t} \prod_{i \in I_s} p_i \qquad (p_i = \Pr_{\mu}[E_i]),$$

$$\mathbb{E}[\#\text{iterations}] \leq \sum_{t=0}^{\infty} \sum_{(I_1,\ldots,I_t) \in Stab} p(I_1, \ldots, I_t).$$

Summation identity: [Kolipaka-Szegedy '11]

If Shearer's conditions are satisfied, then

$$\sum_{t=0}^{\infty} \sum_{(I_1,\ldots,I_t) \in Stab} p(I_1, \ldots, I_t) = \frac{1}{q(p_1, \ldots, p_n)}$$

where $q(p_1, \ldots, p_n) = \sum_{I \in \mathrm{Ind}} (-1)^{|I|} \prod_{i \in I} p_i$.

Analysis with slack

Conditions with slack: Suppose $p'_i = (1+\epsilon)p_i$ and $q_S(p'_1, p'_2, \ldots, p'_n) > 0 \; \forall S \subseteq [n]$. Then

$$\Pr[\#\text{iterations } \geq t] \leq \sum_{(I_1, I_2, \ldots, I_t) \in Stab} p(I_1, \ldots, I_t) \leq \frac{e^{-\epsilon t}}{q(p'_1, \ldots, p'_n)}.$$

We also prove: under an ϵ slack, $q(p'_1, \ldots, p'_n) \geq \epsilon^n$.

Analysis with slack

Conditions with slack: Suppose $p'_i = (1+\epsilon)p_i$ and $q_S(p'_1, p'_2, \ldots, p'_n) > 0 \; \forall S \subseteq [n]$. Then

$$\Pr[\#\text{iterations } \geq t] \leq \sum_{(I_1, I_2, \ldots, I_t) \in Stab} p(I_1, \ldots, I_t) \leq \frac{e^{-\epsilon t}}{q(p'_1, \ldots, p'_n)}.$$

We also prove: under an ϵ slack, $q(p'_1, \ldots, p'_n) \geq \epsilon^n$.

Corollary:

With high prob., the algorithm stops within $O(\frac{n}{\epsilon} \log \frac{1}{\epsilon})$ iterations.

Automatic slack for LLL conditions

Lemma (Harvey-V. '15)

If $p_i \leq x_i \prod_{j \in \Gamma^+(i)} (1 - x_j)$ *then for* $p'_i = \left(1 + \frac{1}{2\sum_{i=1}^{n} \frac{x_i}{1-x_i}}\right) p_i$,

$$q_S(p'_1, \ldots, p'_n) \geq \frac{1}{2} \prod_{i \in S} (1 - x_i).$$

I.e., when the LLL conditions are tight, there is still a slack of

$$\epsilon = \frac{1}{2\sum_{i=1}^{n} \frac{x_i}{1-x_i}}$$

w.r.t. Shearer's conditions.

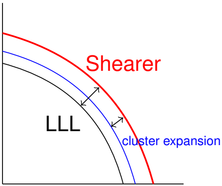

Corollary:
$\mathbb{E}[\#\textit{iterations}] = O((\sum_{i=1}^{n} \frac{x_i}{1-x_i})^2)$ under LLL conditions.

The latest news

[Achlioptas-Iliopoulos '14]:

- do not draw a formal connection with LLL;
- on the other hand claim to go *beyond* the LLL in some ways (orthogonal to Shearer's extension);
- neither framework subsumes the other.

Update: [Achlioptas-Iliopoulos '16], [Kolmogorov '16]

- extended their framework to incorporate resampling operations

Meanwhile, we extended our framework as well...

- both frameworks are becoming one
- unifying concept — *approximate resampling operations*
- this captures exactly the following form of Shearer's Lemma...

Shearer's Lemma with lopsided conditioning

Lemma

Let $E_1, \ldots, E_n$ be events with a graph G such that for every E_i and every event F monotonically depending on $(E_j : j \notin \Gamma^+(i))$,

$$\Pr[E_i \mid F] \leq p_i.$$

Let $q_S(p_1, \ldots, p_n) = \sum_{\text{indep. } I \subseteq S} (-1)^{|I|} \prod_{i \in I} p_i$.
If $\forall S \subseteq [n], q_S(p_1, \ldots, p_n) > 0$, then

$$\Pr[\bigcap_{i=1}^{n} \overline{E_i}] > 0.$$

Open questions

- Is there a deterministic algorithm to find $\omega \in \bigcap_{i=1}^{n} \overline{E_i}$?
- Can we generate a random sample from $\mu|_{\bigcap_{i=1}^{n} \overline{E_i}}$?