Stanford University | **Spring 2016**
Math 233: Non-constructive methods in combinatorics | **Instructor: Jan Vondrák**
Lecture date: April 20, 2016 | **Scribe: Alexander Dunlap**

Lecture 9. Algorithmic LLL continued

Here we continue the analysis of the Moser-Tardos algorithm, extended to the setting of Shearer's Lemma (by Kolipaka and Szegedy).

Lemma 9.1 *If Shearer's conditions are satisfied ($q_I(p_1, \ldots, p_n) > 0$ for every $I \in \mathrm{Ind}(G)$),*

$$\mathbb{E}[\#\text{ of times we resample } \mathcal{E}_i] \leq \frac{q_{\{i\}}}{q_\emptyset}.$$

We prove this lemma in a series of claims.

Claim 9.2 *If $\mathcal{E}_i$ is resampled (at some point) then there is a witness sequence $\mathcal{I} = (\{i\}, \ldots, I_r)$ that occurs in the execution log.*

Proof: We follow the construction from last lecture: Given the execution log containing a resampling $\mathcal{E}_i^{(t)}$, we construct a directed graph rooted at $\mathcal{E}_i^{(t)}$ and from this we construct a stable set sequence $\mathcal{I} = (I_0, I_1, \ldots, I_r)$ such that $I_0 = \{i\}$. By construction, $\mathcal{I}$ occurs in the log. □

Claim 9.3 *If $\mathcal{E}_i$ is resampled q times, then there are q distinct sequences $\mathcal{I} = (\{i\}, \ldots, I_r)$ that occur in the log.*

Proof: By the previous claim, there is a witness sequence starting with $I_0 = \{i\}$ for each resampling of $\mathcal{E}_i$. Also, for two different resamplings $\mathcal{E}_i^{(s)}, \mathcal{E}_i^{(s')}$, if $s < s'$ then $\mathcal{E}_i^{(s)}$ will be included in the directed graph rooted at $\mathcal{E}_i^{(s')}$ and consequently the sequence witnessing $\mathcal{E}_i^{(s')}$ is strictly larger than (properly contains) the sequence witnessing $\mathcal{E}_i^{(s)}$. Thus all the q witness sequences are distinct. □

Claim 9.4 $\mathbb{E}[\#$ *of times we resample* $\mathcal{E}_i] \leq \sum_{r=0}^{\infty} \sum_{\mathcal{I}=(I_0=\{i\},\ldots,I_r)} \prod_{s=0}^{r} p^{I_s}$.

Proof: By the previous claim, we have a distinct stable set sequence starting from $I_0 = \{i\}$ occurring in the log for each resampling of $\mathcal{E}_i$. Therefore, the expected number of resamplings is upper-bounded by the expected number of stable set sequences occurring in the log. For each particular stable set sequence $\mathcal{I} = (I_0, \ldots, I_r)$, the probability that it occurs is at most $\prod_{s=0}^{r} p^{I_s}$ (the main result of last lecture), which gives the bound on the right-hand side. □

To finish the analysis, we aim to upper-bound the summation of the right-hand side of Claim 9.4. First, note that a stable set sequence that occurs in the log can contain only nonempty sets (by construction); we call such sequences *proper stable set sequences*. However, the notion of a stable set sequence allows empty sets; since we have the property $I_{\ell+1} \subseteq \Gamma^+(I_\ell)$, it means that whenever $I_\ell = \emptyset$, we also have $I_{\ell'} = \emptyset$ for all $\ell' > \ell$. A sequence $(I_0, I_1, \ldots, I_\ell, \emptyset, \ldots, \emptyset)$ can be considered equivalent to $(I_0, I_1, \ldots, I_\ell)$, since the quantity $\prod_{s=1}^{r} p^{I_s}$ is the same for both. For counting purposes, it will be convenient to use this equivalence and sum up over all stable set sequences including empty sets.

Definition 9.5 *Let* $\text{Stab}_t(J)$ *be the set of all stable set sequences* $\mathcal{I} = (I_0, \ldots, I_t)$ *of length* $t+1$ *that start with* $I_0 = J$*. Let* $\text{Prop}_t(J)$ *be the subset of* $\text{Stab}_t(J)$ *consisting of sequences that don't contain the empty set.*

Claim 9.6 *Under Shearer's conditions, for every* $J \in \text{Ind}(G)$ *and for all* $t \geq 0$*, we have*

$$\sum_{\mathcal{I}=(I_0,\ldots,I_t)\in\text{Stab}_t(J)} \prod_{s=0}^{t} p^{I_s} \leq \frac{q_J}{q_\emptyset}.$$

Proof: We proceed by induction on t. For the base case, we have $\mathcal{I} = (\{J\})$. We have $q_J = p^J \breve{q}_{V\setminus\Gamma^+(J)} \geq p^J \breve{q}_V = p^J q_\emptyset$, and $q_\emptyset > 0$, which proves $p^J \leq q_J/q_\emptyset$.

For the inductive step, we get

$$\begin{aligned}
\sum_{\mathcal{I}\in\text{Stab}_t(J)} \prod_{s=0}^{t} p^{I_s} &= p^J \sum_{\substack{L\subseteq\Gamma^+(J)\\ L\in\text{Ind}(G)}} \sum_{\mathcal{I}\in\text{Stab}_{t-1}(L)} \prod_{s=1}^{t} p^{I_s} \\
&\leq p^J \sum_{\substack{L\subseteq\Gamma^+(J)\\ L\in\text{Ind}(G)}} \frac{q_L}{q_\emptyset} \\
&= \frac{p^J}{q_\emptyset} \breve{q}_{V\setminus\Gamma^+(J)} \\
&= \frac{q_J}{q_\emptyset}
\end{aligned}$$

where the inequality is the inductive hypothesis and at the end we use the identities $\breve{q}_{V\setminus S} = \sum_{L\subseteq S} q_L$, $q_J = p^J \breve{q}_{V\setminus\Gamma^+(J)}$ (see Lecture 5, (5.4) and (5.5)). □

We remark that equality actually holds in Claim 9.6 if we take the limit as $t \to \infty$, but we do not prove this here.

Corollary 9.7 *Under Shearer's conditions, for every* $J \in \text{Ind}(G)$*,*

$$\sum_{t=0}^{\infty} \sum_{\mathcal{I}=(I_0,\ldots,I_t)\in\text{Prop}_t(J)} \prod_{s=0}^{t} p^{I_s} \leq \frac{q_J}{q_\emptyset}.$$

Proof: By the equivalence discussed above (padding each sequence by empty sets up to length t'), and by Claim 9.6, we have

$$\sum_{t=0}^{t'} \sum_{\mathcal{I}=(I_0,\ldots,I_t)\in\text{Prop}_t(J)} \prod_{s=0}^{t} p^{I_s} = \sum_{\mathcal{I}=(I_0,\ldots,I_{t'})\in\text{Stab}_{t'}(J)} \prod_{s=0}^{t'} p^{I_s} \leq \frac{q_J}{q_\emptyset}.$$

The right-hand side is independent of t'; therefore, as we take the limit $t' \to \infty$, the bound still holds. □

This proves Lemma 9.1, since by Claim 9.4 and Corollary 9.7, the expected number of resamplings of $\mathcal{E}_i$ is bounded by $q_{\{i\}}/q_\emptyset$. This also finishes the main result of the analysis of Kolipaka and Szegedy, showing that the total expected number of resampling steps is upper-bounded by $\sum_{i=1}^{n} \frac{q_{\{i\}}}{q_\emptyset}$, and under the LLL conditions this bound becomes $\sum_{i=1}^{n} \frac{x_i}{1-x_i}$.

Of course it depends on a particular application whether these bounds are polynomial in n or not (although typically they are). We conclude by showing a more concrete bound, under the assumption that Shearer's conditions are satisfied with a certain slack.

Theorem 9.8 *If* $\Pr[\mathcal{E}_i] \leq p_i$*, and* $\varepsilon > 0$ *is such that* $(1+\varepsilon)\mathbf{p}$ *satisfies Shearer's conditions, then*

$$\mathbb{E}[\#\textit{ of resamplings}] \leq \frac{n}{\varepsilon}.$$

Proof: If $(1+\varepsilon)\mathbf{p}$ is in Shearer's region, we know that $q_\emptyset(p_1, \ldots, (1+\varepsilon)p_i, \ldots, p_n) > 0$, by monotonicity of Shearer's region. Differentiating, we have

$$\frac{\partial q_\emptyset}{\partial p_i} = \frac{\partial}{\partial p_i} \sum_{I \in \mathrm{Ind}(G)} (-1)^{|I|} p^I = \sum_{\substack{J \subseteq V \setminus \Gamma^+(i) \\ J \in \mathrm{Ind}(G)}} (-1)^{|J+i|} p^J = -\breve{q}_{V \setminus \Gamma^+(i)}.$$

Thus we have $q_\emptyset(p_1, \ldots, (1+\varepsilon)p_i, \ldots, p_n) = q_\emptyset(\mathbf{p}) + \varepsilon p_i \frac{\partial q_\emptyset}{\partial p_i}|_{\mathbf{p}} = q_\emptyset(\mathbf{p}) - \varepsilon p_i \breve{q}_{V \setminus \Gamma^+(i)} = q_\emptyset(\mathbf{p}) - \varepsilon q_{\{i\}}(\mathbf{p}) > 0$, which implies

$$\frac{q_{\{i\}}(\mathbf{p})}{q_\emptyset(\mathbf{p})} \leq \frac{1}{\varepsilon}.$$

Lemma 9.1 implies the result. □