6.842 Randomness and Computation September 11, 2017

Lecture 3

Lecturer: Ronitt Rubinfeld *Scribe: Albert Cheu*

1 Introduction

In this lecture, we complete the Moser-Tardos algorithm. Specifically, we prove the following:

Theorem 1 *Suppose we have set system $S_1, S_2, S_3 \dots S_m \subset S$ where for any i, $|S_i| = \ell$ and S_i intersects at most d other sets. If $ce(d+1) \leq 2^{\ell-1}$ for constant $c > 1$, then Moser-Tardos finds a coloring of S without a monochromatic S_i in expected time $poly(m, d, |S|)$.*

We will use the "log of execution" and "witness tree" terminology defined in the prior lecture notes. We note that a node of a witness tree can be written as a log element $(time, T) \in \mathbb{Z} \times \{S_1 \dots S_m\}$.

We will make reference to "(fair) coins" which will be the sources of randomness of each event.

2 Proof

2.1 The τ-check and associated facts

We define a procedure that will not be executed but will be useful for analysis, in much the same way witness trees are useful abstractions.

Definition 2 *Suppose tree τ has nodes that each correspond to a log element $(t, S[t]) \in \mathbb{Z} \times \{S_1 \dots S_m\}$. For our purposes, we assume the elements of S are initially uncolored. The τ-*check procedure *is the following:*

1. *Traverse τ reverse BFS order: For each $S[t]$...*
   (a) *Use fair and new coins to color uncolored elements of $S[t]$.*
   (b) *Assign $mono_t$ to 1 if the set is monochromatic, 0 otherwise.*
   (c) *Use fair and new coins to recolor $S[t]$*
2. *Pass if all $mono_t$ are 1, fail otherwise.*

In τ, let M_t be the event that $mono_t = 1$.

Claim 3 *All M_t are independent of one another.*

Proof Consider some $t, S[t]$. M_t is independent of all $M_{t'}$ where $S[t] \cap S[t'] = \emptyset$; the argument below is therefore "interesting" for overlapping sets.

$mono_t = 1$ if and only if all of $S[t]$'s elements are the same color at the beginning of step t.

If some of the elements are uncolored going into step t, the coloring that happens in (a) is independent of any other coloring: the coins are new and independent of any prior flips.

If some of the elements have already been colored, the colors will be from some prior execution of step (c). That uses coins independent of any prior flips.

We showed that $mono_t = 1$ is independent of prior flips. The coloring that (c) will "overwrite" the colors of $S[t]$ so $mono_t =$ is independent of future flips. ■

Claim 4 *When τ has size s, $Pr(\tau$-*check passes*$) = 2^{-s(\ell-1)}$*

Proof This follows immediately from Claim 3 and the fact that $|S_i| = \ell$ for all given S_i. ■

Claim 5 *$Pr(\tau$ is a witness tree$) \leq Pr(\tau$-check passes$)$*

Proof Record onto list C the coin flips from a run of Moser-Tardos that has witness tree τ. Let $(t, S[t])$ be a node in that tree.

Consider a τ-check procedure where instead of flipping new coins, we read off from C in the following manner. In step (a), the coin for an uncolored element u come from the most recent $t' < t$ where $u \in S[t']$. In step (c), we look at the coins used to recolor $S[t]$ at the end of round t of Moser-Tardos.

τ consists of log elements. We log $(t, S[t])$ only when $S[t]$ is monochromatic just before the recoloring at time t. Thus each $mono_i = 1$ and the τ-check passes.

Because the number of coin assignments that make the τ-check pass is at least the number of coin assignments that make τ a witness tree, $Pr(\tau$ is a witness tree$) \leq Pr(\tau$-check passes$)$. ■

2.2 Other facts

Lemma 6 *The number of potential witness trees τ with root S_i and size $|\tau| = s$ is at most $\binom{s(d+1)}{s-1}$*

Proof Note: Taken from scribe notes of Spring '14.

For every given S_j, fix an order of the sets that intersect with S_j including itself. The order is denoted $S_j^1 \dots S_j^{k[j]}$. By assumption, the degree bound is d so $k[j] \leq d + 1$.

Every parent-child pair in τ shares at least one element.

No child is duplicated; the second instance to be added to τ ought to be at a lower level.

Run BFS from the root S_i with queuing order $S_j^1 \dots S_j^{k[j]}$. Let $S_{v(1)} \dots S_{v(s)}$ denote the sequence of sets visited.

Define table T_τ with s rows and $d+1$ columns. The bit in row x and column y will indicate whether $S_{v(x)}^y$ is a child of $S_{v(x)}$ in τ. Because there are $k[v(x)]$ possible children, cells in row x beyond column $y = k[v(x)]$ are set to 0.

There are $s - 1$ edges in a tree with s nodes, so there are $s - 1$ 1-bits. Because no two witness trees have the same table, the number of witness trees is at most the number of such tables. ■

Observation 7 *As a consequence of the witness tree construction rules, no witness tree is duplicated.*

Observation 8 *From the above, a bound on the expected number of distinct witness trees is equal to the expected log length.*

2.3 Bound on expected log length

From the above,

$$\text{Expected log length} = \sum_{S_i} \sum_{\substack{\tau \text{with} \\ \text{root} S_i}} E[\text{no. of times } \tau \text{ is a witness tree}]$$

From Claims 4 and 5,

$$\leq m \sum_{s=1}^{\infty} \sum_{\substack{\tau \text{with} \\ \text{root} S_i \\ |\tau| = s}} 2^{-s(\ell-1)}$$

From Lemma 6,

$$\leq m \sum_{s=1}^{\infty} \binom{s(d+1)}{s-1} 2^{-s(\ell-1)}$$

$$\leq m \sum_{s=1}^{\infty} (e(d+1))^s 2^{-s(\ell-1)}$$

By the assumption in Theorem 1, this is a geometric sequence so

$$= O(m)$$

Each step of the log corresponds to a recoloring which takes $poly(m, d, |S|)$ time. Thus, in expectation, the running time is polynomial.