Principles of Distributed Computing

Lecture 03

Lecturer: Mohsen Ghaffari *Scribe:*

1 Deterministic Coloring of General Graphs

In this section, we start the study of LOCAL coloring algorithms for general graphs. Throughout, the ultimate goal would be to obtain $(\Delta+1)$-coloring of the graphs — that is, an assignment of colors $\{1, 2, \ldots, \Delta+1\}$ to vertices such that no two adjacent vertices receive the same color — where Δ denotes the maximum degree. Notice that by a simple greedy argument, each graph with maximum degree at most Δ has a $(\Delta+1)$-coloring: color vertices one by one, each time picking a color which is not chosen by the already-colored neighbors. However, this greedy argument does not lead to an efficient LOCAL procedure for finding such a coloring. The straightforward transformation of this greedy approach to the LOCAL model would be an algorithm that may need $\Omega(n)$ rounds.

We start with presenting an $O(\log^* n)$-round algorithm that computes a $O(\Delta^2)$ coloring. This algorithm is known as Linial's coloring algorithm [Lin87, Lin92]. Afterward, we discuss how to transform this coloring into a $(\Delta+1)$-coloring.

1.1 Take 1: Linial's Coloring Algorithm

Theorem 1. *There is a deterministic distributed algorithm in the* LOCAL *model that colors any n-node graph G with maximum degree Δ using $O(\Delta^2)$ colors, in $O(\log^* n)$ rounds.*

Outline of the Approach for Theorem 1. The core ingredient of the algorithm is a single-round color reduction method, as we will describe in Section 1.1.1. That will allows us to transform any given coloring with some k colors to some other coloring with a much smaller number $k' \ll k$ of colors. Then, as we discuss in Section 1.1.2, by repeated applications of this single-round color reduction, we obtain the coloring algorithm as claimed in Theorem 1.

1.1.1 Single-Round Color Reduction

Lemma 2. *Given a k-coloring ϕ_{old} of a graph with maximum degree Δ, in a single round, we can compute a k'-coloring ϕ_{new}, for $k' = O(\Delta^2 \log k)$. Furthermore, if $k \leq \Delta^3$, then the bound can be improved to $k' = O(\Delta^2)$.*

The key concept in our single-round color reduction is a combinatorial notion called cover free families, as we will define next.

Definition 3. ***(Cover free families)*** *Given a ground set $\{1, 2, \ldots, k'\}$, a family of sets $S_1, S_2, \ldots, S_k \subseteq \{1, 2, \ldots, k'\}$ is called a Δ-cover free family if for each set of indices $i_0, i_1, i_2, \ldots, i_\Delta \in \{1, 2, \ldots, k\}$, we have $S_{i_0} \setminus \left(\cup_{j=1}^{\Delta} S_{i_j}\right) \neq \emptyset$. That is, if no set in the family is a subset of the union of Δ other sets.*

Using cover free families for color reduction. We use cover free families for color reduction in the obvious way: consider an old coloring ϕ_{old} with k colors and suppose we want a new coloring ϕ_{new} with k' colors. Each node v of old color $\phi_{old}(v) = q$ for $q \in \{1, \ldots, k\}$ will use the set $S_q \subseteq \{1, \ldots, k'\}$ in the cover free family as its *color-set*, i.e., its list of possible colors. Then, it sets its new color $\phi_{new}(v) = q'$ where $q' \in S_q$ is such that q' is not in the color-set of any of the neighbors. Such a color q' is promised to exist, by the definition of cover free families.

As clear from the above outline, we would like to have k' as small as possible, as a function of k and Δ. This would allow us to reduce the number of colors faster. In the following, we prove the existence of Δ-cover free families with a small enough ground set size k'. In particular, Lemma 4 achieves $k' = O(\Delta^2 \log k)$ and Lemma 5 shows that this bound can be improved to $k' = O(\Delta^2)$, if $k \leq \Delta^3$. Toward the end of this subsection, we provide the formal proof that these imply Lemma 2.

Lemma 4. ***(Existence of cover free families)*** *For any k and Δ, there exists a Δ-cover free family of size k on a ground set of size $k' = O(\Delta^2 \log k)$.*

Proof. We use the probabilistic method [AS04] to argue that there exists a Δ-cover free family of size k on a ground set of size $k' = O(\Delta^2 \log k)$. Let $k' = C\Delta^2 \log k$ for a sufficiently large constant $C \geq 2$. For each $i \in \{1, 2, \ldots, k\}$, define each set $S_i \subset \{1, 2, \ldots, k'\}$ randomly by including each element $q \in \{1, 2, \ldots, k'\}$ in S_i with probability $p = 1/\Delta$. We argue that this random construction is indeed a Δ-cover free family, with probability close to 1. Therefore, such a cover free family exists.

First, consider an arbitrary set of indices $i_0, i_1, i_2, \ldots, i_\Delta \in \{1, 2, \ldots, k\}$. We would like to argue that $S_{i_0} \setminus \big(\cup_{j=1}^{\Delta} S_{i_j}\big) \neq \emptyset$. For each element $q \in \{1, 2, \ldots, k'\}$, the probability that $q \in S_{i_0} \setminus \big(\cup_{j=1}^{\Delta} S_{i_j}\big)$ is at exactly $\frac{1}{\Delta}(1-\frac{1}{\Delta})^{\Delta} \geq \frac{1}{4\Delta}$. Hence, the probability that there is no such element q that is in $S_{i_0} \setminus \big(\cup_{j=1}^{\Delta} S_{i_j}\big)$ is at most $(1-\frac{1}{4\Delta})^{k'} \leq exp(-C\Delta \log k/4)$. This is an upper bound on the probability that for a given set of indices $i_0, i_1, i_2, \ldots, i_\Delta \in \{1, 2, \ldots, k\}$, the respective sets violate the cover-freeness property that $S_{i_0} \setminus \big(\cup_{j=1}^{\Delta} S_{i_j}\big) \neq \emptyset$.

There are $k\binom{k-1}{\Delta}$ way to choose such a set of indices $i_0, i_1, i_2, \ldots, i_\Delta \in \{1, 2, \ldots, k\}$, k ways for choosing the central index i_0 and at most $(k-1)^\Delta$ ways for choosing the indices $i_1, i_2, \ldots, i_\Delta$. Hence, by a union bound over all these choices, the probability that the construction fails is at most

$$\begin{aligned} k(k-1)^{\Delta} \cdot exp(-C\Delta \log k/4) &= exp\big(\log k + \Delta(\log(k-1)) - C\Delta \log k/4\big) \\ &\leq exp(-C\Delta \log k/8) \ll 1, \end{aligned}$$

for a sufficiently large constant C. That is, the random construction succeeds to provide us with a valid Δ-cover free family with a positive probability, and in fact with a probability close to 1. Hence, such a Δ-cover free family exists. □

Lemma 5. *For any k and $\Delta \geq k^{1/3}$, there exists a Δ-cover free family of size k on a ground set of size $k' = O(\Delta^2)$.*

Proof. Here, we use an algebraic proof based on low-degree polynomials. Let q be a prime number that is in $[3\Delta, 6\Delta]$. Notice that such a prime number exists by Bertrand's postulate (also known as Bertrand-Chebyshev Theorem). Let $\mathbb{F}_q$ denote the prime field[1] of order q (i.e., integers modulo q). For each $i \in \{1, 2, \ldots, k\}$, associate with set S_i — to be constructed — a distinct degree $d = 2$ polynomial $g_i : \mathbb{F}_q \to \mathbb{F}_q$ over $\mathbb{F}_q$. Notice that there are $q^{d+1} > \Delta^3 \geq k$ such polynomials and hence such an association is possible. Let S_i be the set of all evaluation points of g_i, that is, let $S_i = \{(a, g_i(a)) \,|\, a \in \mathbb{F}_q\}$. These are subsets of the $k' = q^2$ cardinality set $\mathbb{F}_q \times \mathbb{F}_q$. Notice two key properties:

(A) for each $i \in \{1, 2, \ldots, k\}$, we have $|S_i| = q$.

(B) for each $i, i' \in \{1, 2, \ldots, k\}$ such that $i \neq i'$, we have $|S_i \cap S_{i'}| \leq d$.

The latter property holds because, in every intersection point, the degree d polynomial $g_i - g_{i'}$ evaluates to zero, and each degree d polynomial has at most d zeros. Now, the Δ cover-freeness property follows trivially from (A) and (B), because for any set of indices $i_0, i_1, i_2, \ldots, i_\Delta \in \{1, 2, \ldots, k\}$, we have

$$\begin{aligned} |S_{i_0} \setminus \big(\cup_{j=1}^{\Delta} S_{i_j}\big)| \geq |S_{i_0}| &- \sum_{j=1}^{\Delta} |S_{i_0} \cap S_{i_j}| \\ \geq q &- \Delta \cdot d \quad = q - 2\Delta \geq \Delta \geq 1. \end{aligned}$$

□

Remark One can easily generalize the construction of Lemma 5, by taking higher-degree polynomials, to a ground set of size $k' = O(\Delta^2 \log_\Delta^2 k)$, where no assumption on the relation between k and Δ would be needed.

Proof Sketch of Lemma 2. Follows from the existence of cover free families as proven in Lemma 4 and Lemma 5. Namely, each node v of old color $\phi_{old}(v) = q$ for $q \in \{1, \ldots, k\}$ will use the set $S_q \subseteq \{1, \ldots, k'\}$ in the cover free family as its *color-set*. Then, it sets its new color $\phi_{new}(v) = q'$ for a $q' \in S_q$ such that q' is not in the color-set of any of the neighbors. By the definition of the cover free families, and given that ϕ_{old} was a proper coloring, we are guaranteed that such a color q' exists. By the choice of q', the coloring ϕ_{new} is also a proper coloring. □

[1] See https://en.wikipedia.org/wiki/Finite_field

1.1.2 Proving Theorem 1

We now discuss how we obtain Theorem 1 via repeated invocations of Lemma 2.

Proof of Theorem 1. The proof is via iterative applications of Lemma 2. We start with the initial numbering of the vertices as a straightforward n-coloring. With one application of Lemma 2, we transform this into a $O(\Delta^2 \log n)$ coloring. With another application, we get a coloring with $O(\Delta^2(\log \Delta + \log \log n))$ colors. With another application, we get a coloring with $O(\Delta^2(\log \Delta + \log \log \log n))$ colors. After $O(\log^* n)$ applications, we get a coloring with $O(\Delta^2 \log \Delta)$ colors[2]. At this point, we use one extra iteration, based on the second part of Lemma 2, which gets us to an $O(\Delta^2)$-coloring. □

1.2 Take 2: Kuhn-Wattenhofer Color Reduction Algorithm

In the previous section, we saw an $O(\log^* n)$-round algorithm for computing a $O(\Delta^2)$-coloring. In this section, we explain how to transform this into a $(\Delta + 1)$-coloring. We will first see in Section 1.2.1 a basic algorithm that performs this transformation in $O(\Delta^2)$ rounds. Then, in Section 1.2.2, we see how with the addition of a small but clever idea of [KW06], this transformation can be performed in $O(\Delta \log \Delta)$ rounds. As the end result, we get an $O(\Delta \log \Delta + \log^* n)$-round algorithm for computing a $(\Delta + 1)$-coloring.

1.2.1 Warm up: One-By-One Color Reduction

Lemma 6. *Given a k-coloring ϕ_{old} of a graph with maximum degree Δ where $k \geq \Delta + 2$, in a single round, we can compute a $(k-1)$-coloring ϕ_{new}.*

Proof. For each node v such that $\phi_{old}(v) \neq k$, set $\phi_{new}(v) = \phi_{old}(v)$. For each node v such that $\phi_{old}(v) = k$, let node v set its new color $\phi_{new}(v)$ to be a color $q \in \{1, 2, \ldots, \Delta + 1\}$ such that q is not taken by any of the neighbors of u. Such a color q exists, because v has at most Δ neighbors. The resulting new coloring ϕ_{new} is a proper coloring. □

Theorem 7. *There is a deterministic distributed algorithm in the* LOCAL *model that colors any n-node graph G with maximum degree Δ using $\Delta + 1$ colors, in $O(\Delta^2 + \log^* n)$ rounds.*

Proof. First, compute an $O(\Delta^2)$-coloring in $O(\log^* n)$ rounds using the algorithm of Theorem 1. Then, apply the one-by-one color reduction of Lemma 6 for $O(\Delta^2)$ rounds, until getting to a $(\Delta+1)$-coloring. □

1.2.2 Parallelized Color Reduction

Lemma 8. *Given a k-coloring ϕ_{old} of a graph with maximum degree Δ where $k \geq \Delta+2$, in $O(\Delta \log(\frac{k}{\Delta+1}))$ rounds, we can compute a $(\Delta + 1)$-coloring ϕ_{new}.*

Proof. If $k \leq 2\Delta + 1$, the lemma follows immediately from applying the one-by-one color reduction of Lemma 6 for $k - (\Delta + 1)$ iterations. Suppose that $k \geq 2\Delta + 2$. Bucketize the colors $\{1, 2, \ldots, k\}$ into $\lfloor \frac{k}{2\Delta+2} \rfloor$ buckets, each of size exactly $2\Delta + 2$, except for one last bucket which may have size between $2\Delta + 2$ to $4\Delta + 3$. We can perform color reductions in all buckets in parallel (why?). In particular, using at most $3\Delta + 2$ iterations of one-by-one color reduction of Lemma 6, we can recolor nodes of each bucket using at most $\Delta + 1$ colors. Considering all buckets, we now have at most $(\Delta + 1)\lfloor \frac{k}{2\Delta+2} \rfloor \leq k/2$ colors. Hence, we managed to reduce the number of colors by a 2 factor, in just $O(\Delta)$ rounds. Repeating this procedure for $\lceil \log(\frac{k}{\Delta+1}) \rceil$ iterations gets us to a coloring with $\Delta + 1$ colors. The round complexity of this method is $O(\Delta \log(\frac{k}{\Delta+1}))$, because we have $\lceil \log(\frac{k}{\Delta+1}) \rceil$ iterations and each iteration takes $O(\Delta)$ rounds. □

Theorem 9. *There is a deterministic distributed algorithm in the* LOCAL *model that colors any n-node graph G with maximum degree Δ using $\Delta + 1$ colors, in $O(\Delta \log \Delta + \log^* n)$ rounds.*

Proof. First, compute an $O(\Delta^2)$-coloring in $O(\log^* n)$ rounds using the algorithm of Theorem 1. Then, apply the parallelized color reduction of Lemma 8 to transform this into a $(\Delta+1)$-coloring, in $O(\Delta \log \Delta)$ additional rounds. □

[2] If the related calculations are not clear, please ask during the exercise sessions.

References

[AS04] Noga Alon and Joel H Spencer. *The probabilistic method.* John Wiley & Sons, 2004.

[KW06] Fabian Kuhn and Rogert Wattenhofer. On the complexity of distributed graph coloring. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, pages 7–15. ACM, 2006.

[Lin87] Nathan Linial. Distributive graph algorithms global solutions from local data. In *Proc. of the Symp. on Found. of Comp. Sci. (FOCS)*, pages 331–335. IEEE, 1987.

[Lin92] Nathan Linial. Locality in distributed graph algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992.