

Lecture 5: Iterated Rounding

Recall the LP-methodology for designing approximation algorithms. Write a linear integer program capturing the optimum of the problem, solve the linear programming relaxation, and round the possibly fractional solution of the LP to an integer solution. In the last two lectures, we solved the linear program once and spent most of our energy rounding the solution thus obtained into a fully integral solution. In iterated rounding, as the name suggests, the rounding of fractional variables is done in iterations. A quick description of the schema is as follows. In each iteration, a set of variables are made integral; this leads rise to a *residual* problem, and more often than not, the residual problem is an instance of the same class of problems. One then considers the LP relaxation on this residual instance, and if one can argue that the drop in the LP-value is at least $\frac{1}{\rho}$ times the cost of the rounded variables (in case of minimization problems) or at most $\frac{1}{\rho}$ (in case of maximization problems), then one gets a ρ -approximation. Let us elaborate.

Suppose we have a minimization problem, and suppose the LP relaxation for the problem is

$$\min\{c \cdot x : Ax \geq b, x \in [0, 1]^m\}$$

We assume A, b, c are all non-negative. Let x^* be an optimal solution of the above. If $x_k^* = 0$, then in this iteration set the k th variable to 0 and remove it from the problem. The integral cost increases by 0, and it's not hard to see the residual LP is at most the original LP. More interestingly, suppose $x_k^* \geq \frac{1}{\rho}$ for $\rho \geq 1$. In this iteration, round x_k^* to 1 and consider the following LP relaxation for the residual problem.

$$\min\left\{\sum_{j \neq k} c_j x_j : \forall i, \sum_{j \neq k} A_{ij} x_j \geq b_i - A_{ik}, x \in [0, 1]^m\right\}$$

There are two things we need to check here. First is feasibility: is an integral solution to the above residual program *along* with the integral variables set so far constitute a feasible solution for the whole problem? One can check it does. Second is cost analysis. The increase in the cost of the integral solution in this iteration is c_k . We claim that the drop in the LP solution from the first to the second LP is at least $\frac{c_k}{\rho}$. To see this note that the solution $\{x_j^*\}_{j \neq k}$ is a feasible solution to the latter of cost at most the cost of the former LP minus $\frac{c_k}{\rho}$. Now suppose at each iteration we can keep on finding a variable with $x_k^* \geq \frac{1}{\rho}$. Then we can keep repeating this, and finally, when all the variables have been rounded the cost of the LP is zero. Thus the total drop in LP value is the value of the original LP. Since the total integral cost is at most ρ times the total drop, we get a ρ -factor approximation.

The maximization case is trickier. If some $x_k^* \geq 1/\rho$, then we can't simply round it to 1 and argue for the residual linear program similar to what described above. This is because the remaining

solution $\{x_j^*\}_{j \neq k}$ isn't a feasible solution to $\{\forall i, \sum_{j \neq k} A_{ij}x_j \leq b_i - A_{ik}\}$. Nevertheless, iterated rounding is used for maximization problems by setting a set of carefully chosen variables to 0 when a variable is set to 1.

So, at least for minimization problems, it suffices to show that one can always find an optimal solution to the LP relaxation with *at least one* variable having value at least $\frac{1}{\alpha}$, to give an α -approximation. This is done by arguing about a basic feasible solution of the LP; recall there is always a basic feasible optimal solution.

Matchings in Bipartite Graphs

Consider the min-cost perfect matching problem in bipartite graphs. Given a bipartite graph $G = (A, B, E)$ with costs on edges, we need to find a perfect matching of minimum cost. An LP relaxation for the problem is as follows

$$\begin{array}{ll} \min & \sum_{e \in E} c_e x_e \\ \text{subject to} & x(\delta(a)) = 1 \quad \forall a \in A \\ & x(\delta(b)) = 1 \quad \forall b \in B \\ & 0 \leq x_e \leq 1 \quad \forall e \in E \end{array}$$

Let x be a basic feasible solution to the above LP.

Lemma 1. *There is an edge e with $x_e^* = 0$ or $x_e^* = 1$.*

Proof. Suppose not and this all edges e have $0 < x_e^* < 1$. The constraints of the LP imply that each vertex in A (or B) must have at least two edges incident on it; so $|E| \geq 2n$, where $|A| = |B| = n$. Since x^* is basic feasible, we know that $|E|$ is at most the number of linearly independent equalities. These equalities must correspond to the $2n$ constraints shown above because none of the x_e^* 's are 0 or 1. However these are not linearly independent – the sum of the rows corresponding to vertices in A equal the sum of rows corresponding to B . This is a contradiction, implying some x_e^* is either 0 or 1. \square

If $x_e^* = 0$, we delete e from the graph. If $x_e^* = 1$ then we pick this in our matching and delete the end points of e from the graph. In both cases we end up with another min-cost perfect matching problem, and hence we can apply iterated rounding to obtain an exact algorithm.

Survivable Network Design

In this problem, we are given an undirected graph $G = (V, E)$, source-sink pairs $\{s_i, t_i\}$ with requirements r_i , and costs on edges. The goal is to find the cheapest subgraph of G such that there exists r_i edge disjoint paths from s_i to t_i . Kamal Jain, in a breakthrough work, gave a 2-approximation for the problem which we now describe.

Given a subset S of vertices, let $\delta(S)$ denote the set of edges in the cut (S, \bar{S}) and let

$$r(S) := \max_{i: |S \cap \{s_i, t_i\}| = 1} r_i$$

be the minimum number of edges in $\delta(S)$ we need in the solution. An LP relaxation for SNDP is then as follows.

$$\min \quad \sum_{e \in E} c_e x_e \tag{1}$$

$$\text{subject to} \quad x(\delta(S)) \geq r(S) \quad \forall S \subseteq V \tag{2}$$

$$0 \leq x_e \leq 1 \quad \forall e \in E \tag{3}$$

Using the ellipsoid method, one can find a basic feasible optimal solution to the above LP. The separation oracle checks whether a flow of r_i units can be sent from s_i to t_i when each edge has capacity x_e ; this is a max-flow computation. To get a basic feasible solution is more involved; refer the notes for Lecture 3 for a possible approach.

Let x be a basic feasible solution to (1). If $x_e = 0$ for some edge e , delete that edge and resolve the residual LP. (Actually this is not needed, as removing the $x_e = 0$ edges results in a basic feasible solution of the residual LP). Suppose $x_e > 0$ for all $e \in E$. The following is a beautiful theorem of Kamal Jain.

Theorem 1. *There exists an edge e with $x_e \geq 1/2$.*

The basic idea of the proof is as in the matching example above: we suppose the contrary, that is all $x_e < 1/2$, use this to deduce a lower bound on the size of the fractional support of x (the edges e with $1 > x_e > 0$), and show a contradiction by bounding the number of linearly independent constraints of type (2) which hold with equality. Let's start with the last part.

Linearly independent tight constraints of (1)

We start off by noting the following property of the function $r(S)$ which is called *weak supermodularity*.

Claim 1. *For any two subsets $A, B \subseteq V$, one of the two holds*

1. $r(A \cup B) + r(A \cap B) \geq r(A) + r(B)$
2. $r(A \setminus B) + r(B \setminus A) \geq r(A) + r(B)$.

Proof. (Exercise) □

Given a subset $S \subseteq V$, let χ_S denote the row in the constraint matrix corresponding to set S .

Claim 2. *In an undirected graph G , if A, B are two subsets of vertices and $x_e > 0$ is the weight on edge e , then*

1. $x(\delta(A \cup B)) + x(\delta(A \cap B)) \leq x(\delta(A)) + x(\delta(B))$, and
2. $x(\delta(A \setminus B)) + x(\delta(B \setminus A)) \leq x(\delta(A)) + x(\delta(B))$.

Furthermore, equality holds iff $\chi_A + \chi_B = \chi_{A \cup B} + \chi_{A \cap B}$ for case 1 and $\chi_A + \chi_B = \chi_{A \setminus B} + \chi_{B \setminus A}$ for case 2.

Proof. (Exercise) □

Call a set S tight if $x(\delta(S)) = r(S)$. We abuse notation and call a collection of tight sets (S_1, \dots, S_k) linearly independent if $(\chi_{S_1}, \dots, \chi_{S_k})$ are linearly independent. We now bound the size of a maximal collection of linearly independent tight sets.

Claim 3. *If A and B are tight, then either (a) $A \cup B, A \cap B$ are tight, or (b) $A \setminus B, B \setminus A$ are tight. Furthermore, in case (a), we have $\chi_A + \chi_B = \chi_{A \cup B} + \chi_{A \cap B}$, and in case (b) we have $\chi_A + \chi_B = \chi_{A \setminus B} + \chi_{B \setminus A}$. Furthermore,*

Proof. From Claim 1, we have either (a') $r(A \cap B) + r(A \cup B) \geq r(A) + r(B)$, or (b') $r(A \setminus B) + r(B \setminus A) \geq r(A) + r(B)$. If (a') is true, then we get

$$r(A \cap B) + r(A \cup B) \geq r(A) + r(B) = x(\delta(A)) + x(\delta(B)) \geq x(\delta(A \cup B)) + x(\delta(A \cap B)) \geq r(A \cup B) + r(A \cap B)$$

The first equality follows from tightness of A and B , the second from Claim 2, and the last inequality follows from feasibility of x . Since there is equality everywhere, we get $A \cup B$ and $A \cap B$ are tight and $\chi_A + \chi_B = \chi_{A \cup B} + \chi_{A \cap B}$. This is part (a) of the claim.

If (b') is true, we get case (b) of the claim. □

The above claim lets us describe the structure of a maximal collection of linearly independent tight sets. Note there could be many such maximal collections - however, they all are of the same cardinality which equals the rank of the equality matrix.

Definition 1. *A collection of sets \mathcal{L} is called laminar if for any two sets A and B in \mathcal{L} , either A and B are disjoint, or one is a subset of another.*

Let \mathcal{T} be the collection of all tight sets. Pick \mathcal{L} to be a maximal laminar subcollection of \mathcal{T} , that is, any set in $\mathcal{T} \setminus \mathcal{L}$ intersects with some set in \mathcal{L} and doesn't contain or is contained by it. Abusing notation, let $\text{span}(\mathcal{T})$ denote the linear span of the rows $\{\chi_S : S \in \mathcal{T}\}$. Similarly define $\text{span}(\mathcal{L})$. The following claim shows that \mathcal{L} contains a basis for \mathcal{T} and thus *contains* a maximal collection of linearly independent tight set.

Claim 4. $\text{span}(\mathcal{L}) = \text{span}(\mathcal{T})$.

Proof. Since $\mathcal{L} \subset \mathcal{T}$, $\text{span}(\mathcal{L}) \subseteq \text{span}(\mathcal{T})$. Suppose there is a tight set $S \in \mathcal{T} \setminus \mathcal{L}$ such that $\chi_S \notin \text{span}(\mathcal{L})$. Since $S \notin \mathcal{L}$, there are sets in \mathcal{L} which intersect it but don't contain/are contained by it. Choose S so that the number of such sets in \mathcal{L} is minimized, and $L \in \mathcal{L}$ be a set S intersects. By Claim 3, we have one of two cases. Case (a): $S \cup L$ and $S \cap L$ are tight. If either of them is not in $\text{span}(\mathcal{L})$, then that set intersects with strictly smaller number of sets in \mathcal{L} . In particular, it doesn't intersect with L . By our choice of S , both $S \cap L$ and $S \cup L$ are in $\text{span}(\mathcal{L})$. But $\chi_S = \chi_L - \chi_{S \cap L} - \chi_{S \cup L}$ which contradicts $S \notin \text{span}(\mathcal{L})$. Case (b): $S \setminus L$ and $L \setminus S$ are tight, and the argument is similar. □

So we know that the number of linearly independent equalities of (1) is at most the size of a laminar family \mathcal{L} . The next claim shows that this is at most $(2n - 1)$.

Claim 5. *Given any laminar family \mathcal{L} consisting of distinct subsets of a universe U of size n , then $|\mathcal{L}| \leq (2n - 1)$.*

Proof. Remove all singleton sets from \mathcal{L} ; there are at most n such sets. We show that the number of non-singleton sets of \mathcal{L} is at most $(n - 1)$ completing the proof.

Construct the following graph $G_{\mathcal{L}}$: the nodes are sets in \mathcal{L} and elements in U , and there is an arc from node S to node T iff $T \subset S$ and there is no $U \in \mathcal{L}$ such that $T \subset U \subset S$. There is an arc from S to node $v \in U$ if S is the minimal by inclusion set that contains v . Laminarity ensures the following: (a) in-degree of a node is at most 1. If T (or v) had arcs coming in from S and U , then $T \subset S$ and $T \subset U$. Thus S and U intersect implying $U \subset S$ (or vice-versa). This contradicts the presence of the arc from S to T . Thus $G_{\mathcal{L}}$ is a collection of rooted out-trees. (b) Leaves of each tree are disjoint since they are minimal subsets and hence disjoint.

Also note that the out-degree of each internal node is at least 2 since the sets are distinct. Suppose k leaves of $G_{\mathcal{L}}$ are sets in \mathcal{L} . Then the number of elements in U which are leaves is at most $(n - 2k)$. The total number of leaves are thus at most $(n - k)$. The number of internal nodes is thus at most $(n - k - 1)$. This implies the size of non-singleton elements of \mathcal{L} is at most $k + (n - k - 1) = n - 1$. \square

Proof of Theorem 1

The discussion of the previous subsection tells us that given a basic feasible solution x , the collection of linearly independent equalities of (1) correspond to a laminar family \mathcal{L} of tight sets. If we have all edges e have $0 < x_e < 1/2$, we immediately get that $|E| \leq |\mathcal{L}|$. We now wish to reach a contradiction. This will follow by a rather delicate charging argument.

The idea is as follows. We allocate one unit of charge to each edge $e \in E$. The charge transfer scheme moves charges from edges to charges on sets of \mathcal{L} . At the end of the scheme we ensure that each set in \mathcal{L} has at least one unit of charge and there exists an edge with positive charge on it. Since the total charge doesn't decrease, we get $|E| > |\mathcal{L}|$, a contradiction.

The original charging of Jain was a bit complicated; the following is a simpler scheme due to Nagarajan, Ravi and Singh. For each edge $e = (u, v)$ identify three sets. S_u be the minimal set in \mathcal{L} which contains u and not v . S_v is defined similarly. S_{uv} is the minimal set in \mathcal{L} which contains both u and v . Each such edge transfers x_e amounts of charge to both S_u and S_v and $(1 - 2x_e)$ charge to S_{uv} . If any of these three sets are not defined, that amount of charge remains with e .

Claim 6. *After processing each edge $e \in L$, the charge on each set $S \in \mathcal{L}$ is at least 1.*

Proof. Pick a set $S \in \mathcal{L}$. Let T_1, \dots, T_k be sets in \mathcal{L} which are subsets of S but not of any other set in \mathcal{L} . We know that $x(\delta(S)) = r(S)$, and $x(\delta(T_i)) = r(T_i)$ for all i . Let's divide the edges in $F := \delta(S) \cup \bigcup_{i=1}^k \delta(T_i)$ into the following classes

$$\begin{aligned} F_1 &= \{e \in F : e \in \delta(S), e \in \delta(T_i) \text{ for some } i\} \\ F_2 &= \{e \in F : e \in \delta(S), e \notin \delta(T_i) \text{ for any } i\} \\ F_3 &= \{e \in F : e \in \delta(T_i), e \in \delta(T_j) \text{ for some pair } i, j\} \\ F_4 &= \{e \in F : e \in \delta(T_i), e \notin \delta(T_j) \text{ for any other } j, e \notin \delta(S)\} \end{aligned}$$

What we get from the tightness of all these sets is

$$r(S) = x(\delta(S)) = x(F_1) + x(F_2)$$

and

$$\sum_{i=1}^k r(T_i) = \sum_{i=1}^k x(\delta(T_i)) = x(F_1) + x(F_4) + 2x(F_3)$$

Furthermore, since S, T_1, \dots, T_k are linearly independent, we get that one of F_2, F_3 and F_4 are non-empty. What's the charge assigned to S ? Each edge in F_2, F_3 and F_4 give charge to S : x_e , $(1 - 2x_e)$ and $x_e + (1 - 2x_e) = 1 - x_e$, respectively. Since not all of them are empty, S gets some positive charge. The total charge on S is

$$x(F_1) + |F_3| - 2x(F_3) + |F_4| - x(F_4)$$

Rearranging terms, this is equal to $r(S) - \sum_{i=1}^k r(T_i) + |F_3| + |F_4|$, which is an integer since $r(\cdot)$ is an integer. Since the charge is positive and an integer it must be at least 1. This proves the claim and therefore Theorem 1. \square

Armed with Theorem 1, the iterated rounding algorithm for SNDP is clear. We pick edge e with $x_e \geq 1/2$ and round it to 1. This leads to a residual problem where $r(S)$ of each set containing $e \in \delta(S)$ drops by 1. And then we iterate. Hold on. In our proof of there being an edge e with $x_e \geq 1/2$, we crucially used weak supermodularity of $r(S)$. Do they still hold? If they do, then we are fine and our argument can be iterated. I leave it as an exercise to show that indeed we are fine.

Maximum Matchings in Hypergraphs

(We didn't do this in the lecture, but here are the notes for it anyway.) We end this lecture with an application of iterated rounding to a maximization problem. We'll see how one moves to a residual problem while maintaining that the residual LP has a valid solution which is not too small. The problem is of finding a maximum matching in a hypergraph. A hypergraph $H = (V, E)$ is a generalization of a graph where each edge e , called hyperedge, is a general subset of V . A hypergraph is k -uniform if each edge e has $|e| = k$. A matching is a collection of disjoint hyperedges. Given a k -uniform hypergraph, we wish to find a maximum cardinality matching.

The LP relaxation.

$$\max \sum_{e \in E} x_e \tag{4}$$

$$\text{subject to } \sum_{e: v \in e} x_e \leq 1 \quad \forall v \in V \tag{5}$$

$$0 \leq x_e \leq 1 \quad \forall e \in E \tag{6}$$

Let x be a basic feasible solution to the above LP. We have the following claim.

Claim 7. *There exists an edge $e \in E$ with $x_e \geq 1/k$.*

Proof. Let $F := \{e : x_e > 0\}$. Suppose $x_e < 1$ for all edges, otherwise we are done. Let $U \subseteq V$ be the set of tight vertices: $\sum_{e:v \in e} x_e = 1$. Since x is a bfs, we have $|F| \leq |U|$. So,

$$|U| = \sum_{v \in U} 1 = \sum_{v \in U} \sum_{e:v \in e} x_e = \sum_{e \in F} x_e \sum_{v \in e} 1 = k \sum_{e \in F} x_e$$

If all $x_e < 1/k$, then the final term will be $< |F|$ contradicting $|F| \leq |U|$. \square

Now what do we do? Suppose we pick the edge e^* with $x_{e^*} > 1/k$. Well since in the end we must return a matching, we must set $x_e = 0$ for all e which intersects e^* . Let this set be $E^* := \{e : e \cap e^* \neq \emptyset\}$. In the residual problem, we delete all edges in E^* and vertices in e^* – we get another hypergraph matching problem and a resulting LP. Can we bound the drop in the LP solution. Let LP_1 be the value of the original LP and LP_2 the value of the residual LP.

Claim 8. $LP_1 - LP_2 \leq (k - 1 + 1/k)$

Proof. Note that x_e for all $e \in E \setminus E^*$ is a valid solution to the residual LP. Thus the drop in the LP value is at most $x(E^*)$. Now,

$$x(E^*) = x_{e^*} + \sum_{e \neq e^*, e \in E^*} x_e \leq x_{e^*} + \sum_{v \in e^*} \sum_{e \neq e^*: v \in e} x_e$$

Since $x_{e^*} \geq 1/k$, the feasibility of x gives us $\sum_{e \neq e^*: v \in e} x_e \leq 1 - x_{e^*}$. Thus, we get $x(E^*) \leq x_{e^*} + k(1 - x_{e^*}) = k - (k - 1)x_{e^*} \leq k - \frac{k-1}{k}$ since $x_{e^*} \geq 1/k$. \square

So, the LP drops by at most $(k - 1 + 1/k)$ and the algorithm gains one edge. Since we delete all edges incident on the edge picked, in the end we return a feasible matching which has value at least $\frac{1}{(k-1+1/k)}$ times the original LP value.