

CS264: Beyond Worst-Case Analysis

Lecture #7: Exact Recovery in Perturbation-Stable Minimum Cut Instances*

Tim Roughgarden[†]

January 31, 2017

1 Exact Recovery

This lecture makes explicit a theme introduced surreptitiously last lecture: *exact recovery*. The genre of question in this sequence of lectures is: for an algorithm A of interest, for which inputs does A solve the problem exactly? For example, last lecture, we proved that the single-link++ algorithm recovers the optimal solution in every 2-stable k -median instance.¹

We generally think of the algorithm A above as either: (i) an already-implemented algorithm that is available to us, such as a linear programming solver; or (ii) a relatively simple algorithm that we have a good understanding of and could code up easily. The reason is that, in practice, designing or implementing a new algorithm from scratch is often a last resort. It is usually far preferable if something in the current algorithmic toolbox is good enough, at least for the problem instances relevant to the application. The exact recovery results in this sequence of lectures develop theoretical guarantees that correspond well to this style of thinking.²

This lecture covers another case study on exact recovery, for a graph cut problem. We'll retain from last lecture the definition of γ -perturbation-stable instances and the flavor of the problems studied, as cut problems can be thought of as a type of clustering problem. What will be different is the algorithm studied — linear programming rather than a more elementary algorithm.

*©2014–2017, Tim Roughgarden.

[†]Department of Computer Science, Stanford University, 474 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

¹*Approximate* notions of recovery (under weaker input assumptions) are also very interesting. For simplicity we'll stick with exact recovery for these lectures.

²This contrasts with a majority (but certainly not all) theoretical computer science research on algorithms, where the primary goal is usually to come up with new algorithms that are in some sense “better” than the state-of-the-art (often according to some worst-case measure).

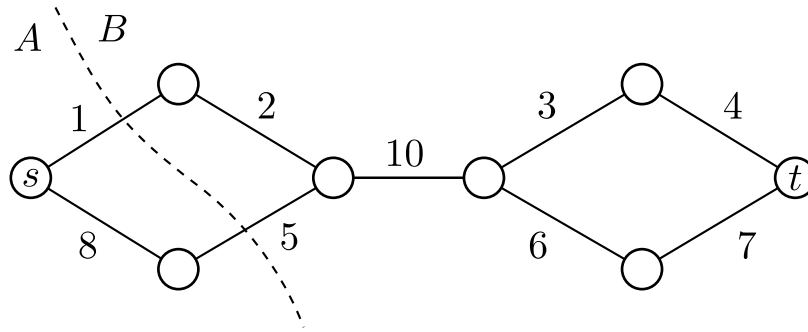


Figure 1: The minimum s - t cut of a graph.

The broader theme, to be elaborated on in future lectures, is to explain when NP -hard problems can be solved exactly via mathematical programming (linear, semidefinite, convex, etc.). For mathematical programs that can be solved in polynomial time, like linear or convex programs, we do not expect exact recovery in worst-case instances. Thus, even more so than in the last lecture, we'll use the stability assumption only in the analysis of an algorithm (linear programming), and not as a guide for designing new algorithms.

2 The Minimum s - t Cut Problem

We begin with a warm-up, the minimum s - t cut problem. Everything in this section has been known for a half-century, but we'll phrase it in a way that generalizes surprisingly effortlessly to NP -hard generalizations of the problem (under stronger assumptions).

Recall the input to the problem is an undirected graph $G = (V, E)$ with a source vertex $s \in V$ and a sink vertex $t \in V$. Every edge $e \in E$ has a positive *cost* $c_e > 0$. The desired output is an s - t cut — a partition of the vertex set V into A and B with $s \in A$ and $t \in B$ — that minimizes the sum of the costs of the cut edges. (An edge is cut if it has one endpoint in each of A and B .) See Figure 1.

As you may know, the minimum s - t cut problem can be solved in polynomial time, for example using a maximum flow algorithm. The point of this section is to show that, alternatively, one can just solve a suitable linear program and the minimum s - t cut pops out as the optimal solution. For simplicity, we assume throughout this section that the optimal solution is unique. When we study harder problems later, we'll upgrade this uniqueness assumption to a suitable γ -perturbation-stable-stability assumption.

Recall that a linear program has a set of real-valued decision variables, and the goal is to optimize an objective function that is linear in the decision variables, subject to constraints that are also linear. Geometrically, the constraints are an intersection of halfspaces, the objective function gives a direction, and the goal is to find the feasible point that is furthest in the desired direction (Figure 2).

How do we encode the minimum s - t cut problem as a linear program? There are two

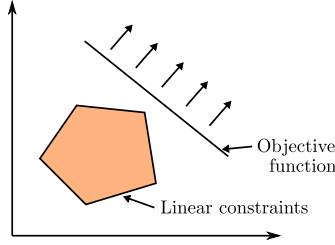


Figure 2: Cartoon representation of a linear program.

sets of decision variables. First, a nonnegative variable $x_e \geq 0$ for each edge $e \in E$, with the semantics that $x_e = 1$ if e gets cut and $x_e = 0$ otherwise. Second, there is a nonnegative decision variable $d_v \geq 0$ for each vertex $v \in V$; the intention is that $d_v = 0$ if v is in the source-side A of the cut, and $d_v = 1$ if $v \in B$. Given these semantics, two obvious constraints are $d_s = 0$ and $d_t = 1$. The more interesting constraints are meant to enforce the property that, whenever the endpoints u, v of an edge e are on different sides of the s - t cut, edge e is cut. Precisely, these constraints are:

$$\begin{aligned} x_e &\geq d_u - d_v \\ x_e &\geq d_v - d_u \end{aligned}$$

for every edge $e = (u, v) \in E$. These constraints are equivalent to the single constraint $x_e \geq |d_u - d_v|$; we have written them so that it is obvious that they are linear constraints. Given the intended semantics for the x -variables, it is clear that the (linear) objective function is

$$\min \sum_{e \in E} c_e x_e.$$

We denote this linear program by (LP1).

It is clear that every s - t cut (A, B) induces a feasible 0-1 solution to (LP1) with matching objective function value: set $d_v = 0$ for $v \in A$, $d_v = 1$ for $v \in T$, and $x_e = 1$ for the cut edges and $x_e = 0$ for the rest. In this sense, (LP1) is a *relaxation* of the s - t minimum cut problem. There are, however, feasible solutions to the linear program that do not correspond to s - t cuts. This is intuitively obvious from a glance at Figure 2 — the s - t cuts correspond to a discrete and finite set of points, whereas the feasible region of a linear program is a convex set. More concretely, in the graph of Figure 1, one fractional solution is to set $x_e = \frac{1}{5}$ for every edge, with d_v -values of vertices increasing by $\frac{1}{5}$ with each hop away from s and toward t .

An important fact, that you should definitely know, is that linear programs can be solved efficiently, both in theory and in practice.

Fact 2.1 *There is a polynomial-time algorithm for linear programming.*

Whether proving theorems or solving real-world problems, most users of linear programming algorithms treat them as “black boxes” and don’t understand how or why they work. We will do the same in the course. If you want to see what’s “under the hood,” check out e.g. [6].

In this section, we prove the following.

Proposition 2.2 *For every minimum s - t cut problem with a unique optimal solution, the optimal solution to the linear program (LP1) is the (straightforward encoding of) the minimum s - t cut.*

The point of Proposition 2.2 is that, while the feasible region of the linear program includes plenty of points that don’t correspond to s - t cuts, we don’t need to worry about them — as long as we solve the linear program to optimality, we’ll get back an integer solution that corresponds to the minimum cut.³

To prove the proposition, let C^* denote the minimum s - t cut, and let F^* denote the edges cut by it. For another s - t cut C , cutting the edges $F(C)$, we let $\Delta(C)$ denote the extent to which the total cost of C exceeds that of C^* :

$$\Delta(C) = \sum_{e \in F(C) \setminus F^*} c_e - \sum_{e \in F^* \setminus F(C)} c_e. \quad (1)$$

Since C^* is the unique minimum s - t cut,

$$\Delta(C) \geq 0 \quad (2)$$

for every s - t cut C , with equality holding if and only if $C = C^*$.

Similarly, for the optimal solution $(\hat{\mathbf{x}}, \hat{\mathbf{d}})$ to (LP1), we denote by $\Delta(\hat{\mathbf{x}})$ the extent to which the objective function value of the integer solution corresponding to C^* exceeds that of $\hat{\mathbf{x}}$:

$$\Delta(\hat{\mathbf{x}}) = \sum_{e \notin F^*} c_e \hat{x}_e - \sum_{e \in F^*} c_e (1 - \hat{x}_e).$$

Since C^* corresponds to a feasible solution to (LP1), and $\hat{\mathbf{x}}$ corresponds to the optimal solution,

$$\Delta(\hat{\mathbf{x}}) \leq 0. \quad (3)$$

The key to proving Proposition 2.2 is the following lemma.

Lemma 2.3 *There is a randomized algorithm that outputs a (random) s - t cut C such that, for every edge $e \in E$,*

$$\Pr[e \text{ cut by } C] = \hat{x}_e \quad (4)$$

and hence

$$\Pr[e \text{ not cut by } C] = 1 - \hat{x}_e. \quad (5)$$

³This is true even when there is not a unique minimum s - t cut. The reasons, which we won’t prove here, are: (i) all of the extreme points (a.k.a. vertices) of the feasible region are integral and correspond to s - t cuts; and (ii) linear programming algorithms always return an extreme point of the feasible region.

Before proving the lemma, we explain why it implies Proposition 2.2. The intuition is that since (i) the optimal fractional solution $\hat{\mathbf{x}}$ can only be better than the optimal integral solution C^* ; and (ii) Lemma 2.3 gives a distribution over s - t cuts that is as good, on average, as C^* , and so the distribution must be a point mass on the cut C^* , which implies that $\hat{\mathbf{x}}$ is a 0-1 solution.

Formally, let C denote the random s - t cut generated by the randomized algorithm in Lemma 2.3, and consider the expected extent to which its cost exceeds that of C^* :

$$\mathbf{E}[\Delta(C)] = \sum_{e \notin F^*} c_e \cdot \Pr[C \text{ cuts } e] - \sum_{e \in F^*} c_e \cdot \Pr[C \text{ doesn't cut } e] \quad (6)$$

$$= \sum_{e \notin F^*} c_e \hat{x}_e - \sum_{e \in F^*} c_e (1 - \hat{x}_e) \quad (7)$$

$$= \Delta(\hat{\mathbf{x}}) \leq 0, \quad (8)$$

where equation (6) follows from linearity of expectation, equation (7) from (4) and (5) in Lemma 2.3, and (8) follows from (3). Since $\Delta(C)$ is always nonnegative and is strictly positive unless $C = C^*$ (recall (2)), the derivation (6)–(8) implies that the randomized algorithm of Lemma 2.3 outputs the cut C^* with probability 1. Inspection of (4) and (5) shows that this occurs only if $\hat{x}_e \in \{0, 1\}$ for every $e \in E$. This completes the proof of Proposition 2.2.

Proof of Lemma 2.3: Let $(\hat{\mathbf{x}}, \hat{\mathbf{d}})$ be the optimal solution to (LP1). Intuitively, the randomized algorithm blows up a balloon around the source s up to a random size, with everything inside the balloon comprising the source-side of the cut. Formally:

s - t Cut Rounding Algorithm

1. Pick a “radius” $r \in (0, 1)$ uniformly at random.
2. Define $A = \{v \in V : \hat{d}_v \leq r\}$, and $B = V \setminus A$.

Since $\hat{d}_s = 0$ and $\hat{d}_t = 1$, the algorithm’s output is an s - t cut with probability 1.

To verify (4) (and hence (5)), first note that the optimality of $(\hat{\mathbf{x}}, \hat{\mathbf{d}})$ implies that $\hat{x}_e = |\hat{d}_u - \hat{d}_v|$ for every edge $e = (u, v)$. (Otherwise decreasing \hat{x}_e would yield another feasible solution, with smaller objective function value.) We thus have

$$\begin{aligned} \Pr[e \text{ cut by } (A, B)] &= \Pr[r \text{ is between } \hat{d}_u \text{ and } \hat{d}_v] \\ &= |\hat{d}_u - \hat{d}_v| \\ &= \hat{x}_e \end{aligned}$$

for every edge $e = (u, v)$. ■

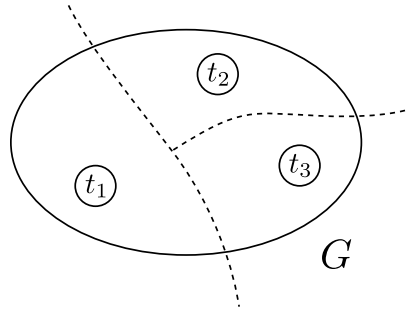


Figure 3: A multiway cut of a graph (with $k = 3$).

3 Perturbation-Stable Instances of the Minimum Multiway Cut Problem

Proposition 2.2 shows that in every instance of the minimum s - t cut problem with a unique optimal solution, the optimal solution can be recovered exactly in polynomial time by solving a linear programming relaxation. This section takes this result to the next level, by considering an NP -hard generalization of the minimum s - t cut problem, replacing the uniqueness assumption with a stronger perturbation-stability assumption, and again showing how to recover the optimal (integral) solution using linear programming.

In the *minimum multiway cut problem*, the input is an undirected graph $G = (V, E)$ with positive edge costs c_e , and a set $t_1, t_2, \dots, t_k \in V$ of *terminals*. A *multiway cut* is a partition of the vertex set V into k groups S_1, \dots, S_k , with $t_i \in S_i$ for $i = 1, 2, \dots, k$. In addition to being a natural generalization of the fundamental s - t minimum cut problem (which is equivalent to the $k = 2$ case), the problem also shows up in applications, such as image segmentation.⁴ The problem is NP -hard for every $k \geq 3$ [3]. See Figure 3.

Since the problem is NP -hard, we do not expect to recover an optimal solution using linear programming in worst-case instances. But we *can* do this under a stability assumption analogous to the one introduced last lecture for k -median instances.

Definition 3.1 For a parameter $\gamma \geq 1$, an instance of multiway cut is γ -*perturbation-stable* if a common partition (S_1^*, \dots, S_k^*) is the unique optimal solution of all γ -perturbations, where each original edge cost c_e is replaced an edge cost $c'_e \in [\frac{1}{\gamma}c_e, c_e]$.

The main result of this lecture is the following result; we'll define the linear programming relaxation (LP2) shortly.

⁴In the most straightforward formulation, vertices correspond to pixels, edges correspond to neighboring pixels, edge costs correspond to a prior guess about how likely two neighboring pixels are to be in the same region (e.g., based on how similar their colors are), and groups of the multiway cut correspond to the different regions of the image [1].

Theorem 3.2 ([5]) *For every 4-perturbation-stable multiway cut instance, the optimal solution to the linear program (LP2) is the (straightforward encoding of) the optimal multiway cut.*

This linear program (LP2) (originally from [2]) is a nice extension of (LP1), which can be viewed as the $k = 2$ special case. We again have a set of decision variables $x_e \geq 0$ for each $e \in E$, meant to indicate which edges get cut. We now have multiple variables $d_v^1, \dots, d_v^k \geq 0$ per vertex, with the idea that if v is assigned to the i th group, then $d_v^i = 1$ and $d_v^j = 0$ for every $j \neq i$. The first set of constraints just asserts that each terminal is fully assigned to its own group:

$$d_{t_i}^i = 1$$

for each $i = 1, 2, \dots, k$. For other vertices, we insist that they are fully assigned, perhaps fractionally, to the groups:

$$\sum_{i=1}^k d_v^i = 1$$

for every $v \in V$. Next, we give constraints stating that edges with endpoints in different groups must be cut. First, we have

$$\begin{aligned} y_e^i &\geq d_u^i - d_v^i \\ y_e^i &\geq d_v^i - d_u^i \end{aligned}$$

for every edge $e = (u, v) \in E$ and $i = 1, 2, \dots, k$, where the y_e^i 's are auxiliary variables. These constraints assert that the difference in the extent to which the edge's endpoints are assigned to group i should contribute to the extent to which e is cut. Finally, for every $e \in E$, we have the constraint

$$x_e = \frac{1}{2} \sum_{i=1}^k y_e^i, \tag{9}$$

which aggregates the group assignment differences across the coordinates. Note that if the endpoints of an edge e are assigned integrally to different groups (i and j , say), then $y_e^i = y_e^j = 1$; this motivates the “ $\frac{1}{2}$ ” in (9), so that $x_e = 1$ for such an edge.⁵ Also note that every multiway cut naturally induces a feasible 0-1 solution to (LP2), with matching objective function value, according to the intended semantics of the decision variables. As another example, if $k = 3$, and the d -values of u and v are $(\frac{1}{2}, \frac{1}{2}, 0)$ and $(0, \frac{1}{2}, \frac{1}{2})$, then the x -value of edge $e = (u, v)$ must be at least $\frac{1}{2}$. Pictorially, we can think of (LP2) as assigning each vertex v a point (d_v^1, \dots, d_v^k) in the k -simplex (Figure 4), with corners of the simplex corresponding to the groups. Constraint (9) states that edge $e = (u, v)$ is cut in proportion to the “distance” (specifically, half of the ℓ_1 distance) between the points of the simplex to which u and v are assigned.

⁵If we were only concerned with integral solutions, the constraint $x_e = \max_{i=1}^k y_e^i$ would already force edge e to be cut whenever its endpoints are assigned to different groups. When the variables can be fractional, however, the constraint (9) rules out many more fractional solutions, thereby increasing the utility of (LP2) as an approximation to the problem that we actually care about, the minimum multiway cut problem.

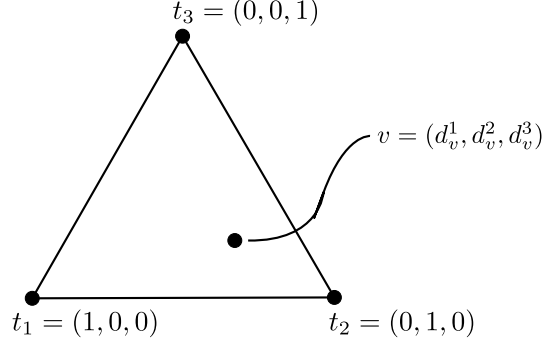


Figure 4: Pictorial representation of (LP2): the simplex when $k = 3$, drawn in two dimensions.

We now turn to Theorem 3.2. The proof follows the same template as that of Proposition 2.2. Let C^* be the unique multiway cut, cutting the edges F^* . For another multiway cut C , cutting edges $F(C)$, define

$$\Delta_4(C) = \frac{1}{4} \sum_{e \in F(C) \setminus F^*} c_e - \sum_{e \in F^* \setminus F(C)} c_e. \quad (10)$$

Recall the interpretation of the quantity $\Delta(C)$ in (1), as the extent to which the $(s-t)$ cut C is more costly than the optimal $(s-t)$ cut C^* . The interpretation of $\Delta_4(C)$ in (10) is as the extent to which the cost of the multiway cut C exceeds that of C^* , *after applying the 4-perturbation most favorable for C* . The 4-perturbation that most narrows the cost gap between the optimal multiway cut C^* and some other multiway cut C is the one that reduces the cost of the edges cut by C but not C^* as much as possible (i.e., by a factor of 4). With this understanding of the expression in (10), it should be clear that $\Delta_4(C^*) = 0$ and $\Delta_4(C) > 0$ for every other multiway cut $C \neq C^*$ (by 4-perturbation-stability).

For the optimal solution $(\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{d}})$ to (LP2), as before we define

$$\Delta(\hat{\mathbf{x}}) = \sum_{e \notin F^*} c_e \hat{x}_e - \sum_{e \in F^*} c_e (1 - \hat{x}_e),$$

as the increase in cost incurred by switching from the (integral) solution corresponding to C^* to the (possibly fractional) optimal LP solution. As before, since C^* induces a feasible solution to (LP2) and $\hat{\mathbf{x}}$ is optimal,

$$\Delta(\hat{\mathbf{x}}) \leq 0. \quad (11)$$

We next give the analog of Lemma 2.3. We can't expect an equally strong lemma for multiway cuts with $k \geq 3$, since this would give a polynomial-time exact algorithm for an NP -hard problem. Instead, we content ourselves with a version in which the identities (4) and (5) hold approximately.

Lemma 3.3 *There is a randomized algorithm that outputs a (random) multiway cut C such that, for every edge $e \in E$,*

$$\Pr[e \text{ cut by } C] \leq 2\hat{x}_e \tag{12}$$

and

$$\Pr[e \text{ not cut by } C] \geq \frac{1 - \hat{x}_e}{2}. \tag{13}$$

Note that (13) does not follow from (12); the implied lower bound of $1 - 2\hat{x}_e$ is not good enough for our purposes.

The proof that Lemma 3.3 implies Theorem 3.2 follows the proof of Proposition 2.2. We lose two factors of 2 in the argument from (12) and (13), but these get absorbed by the 4-perturbation-stability assumption. Formally, let C denote the random multiway cut generated by the randomized algorithm in Lemma 3.3. We have

$$\mathbf{E}[\Delta_4(C)] = \frac{1}{4} \sum_{e \notin F^*} c_e \cdot \Pr[C \text{ cuts } e] - \sum_{e \in F^*} c_e \cdot \Pr[C \text{ doesn't cut } e] \tag{14}$$

$$\leq \frac{1}{4} \sum_{e \notin F^*} c_e (2\hat{x}_e) - \sum_{e \in F^*} c_e \frac{1}{2} (1 - \hat{x}_e) \tag{15}$$

$$= \frac{1}{2} \Delta(\hat{\mathbf{x}}) \tag{16}$$

$$\leq 0,$$

where equation (14) follows from linearity of expectation, inequality (15) from (12) and (13) in Lemma 3.3, and (16) follows from (11). Since $\Delta_4(C)$ is always nonnegative and is strictly positive unless $C = C^*$, the derivation (14)–(16) implies that the randomized algorithm of Lemma 3.3 outputs the cut C^* with probability 1. We'll see below that this occurs only if $\hat{x}_e \in \{0, 1\}$ for every $e \in E$. Modulo this point, this completes the proof of Theorem 3.2.

Proof of Lemma 3.3: Let $(\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{d}})$ be the optimal solution to (LP2). The randomized algorithm (originally from [4]) is simple, though unlike the one in the proof of Lemma 2.3, it's probably not the first one you'd think of:

Multiway Cut Rounding Algorithm

1. Initially, all vertices are unassigned.
2. While there is at least one unassigned vertex:
 - (a) Pick a threshold $r \in (0, 1)$ uniformly at random.
 - (b) Pick a group $i \in \{1, 2, \dots, k\}$ uniformly at random.
 - (c) For every unassigned vertex v , if $\hat{d}_v^i \geq r$, assign vertex v to group i .

The algorithm eventually terminates with a multiway cut — each vertex is assigned to exactly one group, and since $\hat{d}_{t_i}^i = 1$ and $\hat{d}_{t_i}^j = 0$ for $j \neq i$ for every i , every terminal is assigned to

the correct group. We leave the verification of inequalities (12) and (13) to Homework #4.⁶
 ■

Observe that, if \hat{d}_v^j is strictly between 0 and 1 for any vertex v and coordinate j , then $\hat{d}_v^\ell \in (0, 1)$ for some other coordinate ℓ as well (since $\sum_{i=1}^k d_v^i = 1$). In this case, the randomized algorithm of Lemma 3.3 will sometimes assign v to group j , and sometimes to group ℓ . We conclude that the randomized algorithm always returns the same multiway cut only if the optimal fractional solution $(\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{d}})$ is integral. This completes the proof of Theorem 3.2, that in 4-perturbation-stable multiway cut instances, a natural linear programming relaxation (LP2) is guaranteed to be exact.

An interesting open research question is to prove better upper bounds, or any lower bounds, on the minimum value of γ such that the optimal solution of γ -perturbation-stable multiway cut instances can be recovered in polynomial time (by solving (LP2) or by other means).

References

- [1] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, 2011.
- [2] G. Calinescu, H. Karloff, and Y. Rabani. An improved approximation algorithm for multiway cut. *Journal of Computer and System Sciences*, 60(3):564–574, 2000.
- [3] E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis. The complexity of multiterminal cuts. *SIAM Journal on Computing*, 23(4):864–894, 1994.
- [4] J. M. Kleinberg and É. Tardos. Approximation algorithms for classification problems with pairwise relationships: metric labeling and markov random fields. *Journal of the ACM*, 49(5):616–639, 2002.
- [5] K. Makarychev, Y. Makarychev, and A. Vijayaraghavan. Bilu-Linial stable instances of max cut and minimum multiway cut. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 890–906, 2014.

⁶Here’s an example to get a feel for the algorithm’s behavior. Suppose the endpoints u, v of edge e have \hat{d} -values $(\frac{1}{2}, \frac{1}{2}, 0)$ and $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. If i is randomly chosen to be 3, then u is definitely not assigned to any group in this iteration, and v is assigned in this iteration to group 3 with probability $\frac{1}{3}$. In the latter case, edge e will definitely be cut at the end of the day. If i is chosen to be 1, then three things could happen: both u and v could be assigned to group 1 (if $r \leq \frac{1}{3}$), in which case e is definitely not cut; neither could be assigned (if $r > \frac{1}{2}$), leaving the edge’s fate to a future iteration; or only u could be assigned to group 1 (if $r \in (\frac{1}{3}, \frac{1}{2}]$), in which case e is likely to eventually get cut (although there is still a chance v could be assigned to group 1 in a future iteration).

- [6] J. Matousek and B. Gärtner. *Understanding and Using Linear Programming*. Springer, 2006.