Advanced Algorithms 2012A
Lecture 4 – flow/cut gaps*

Robert Krauthgamer

1 Maximum Flow and Minimum Cut (Single Commodity)

Let $G = (V, E)$ be a directed graph with a source $s \in V$, a sink $t \in V$, and edge capacities $c_e \geq 0$. Recall that maximum st-flow is equal to minimum st-cut. We first review material from the previous class, in which Uri gave an alternative proof using LP duality and total unimodularity.

Maximum flow can be written as the following LP with variables $(x_{uv} : uv \in E^*)$. We define $E^* = E \cup \{ts\}$ which adds an edge from t to s.

$$\begin{array}{lll} \text{maximize} & x_{ts} & \\ \text{subject to} & \sum_{v:\ vu \in E^*} x_{vu} - \sum_{v:\ uv \in E^*} x_{uv} = 0 & \forall u \in V \\ & x_{uv} \leq c_{uv} & \forall uv \in E \\ & x_{uv} \geq 0 & \forall uv \in E^* \end{array} \tag{1}$$

The dual LP has variables $(y_u : u \in V)$ and $(y_{uv} : uv \in E)$ and is given by:

$$\begin{array}{lll} \text{minimize} & \sum_{uv \in E} c_{uv} y_{uv} & \\ \text{subject to} & y_u - y_v - y_{uv} \geq 0 & \forall uv \in E \\ & y_t - y_s \geq 1 & \\ & y_{uv} \geq 0 & \forall uv \in E \end{array} \tag{2}$$

This LP is a *relaxation* of the min-cut problem. We can call it "fractional cut". Using also duality, we thus have

max-flow = fractional-cut $\leq$ min-cut

Exer: prove the inequality max-flow $\leq$ min-cut directly by definition.

*These notes summarize the material covered in class, usually skipping proofs, details, examples and so forth, and possibly adding some remarks, or pointers. The exercises are for self-practice and need not be handed in. In the interest of brevity, most references and credits were omitted.

1.1 Using Total Unimodularity

We saw last week how to use TUM to prove that

max-flow = fractional-cut = min-cut.

1.2 Alternative LP formulations

Let P be the collection of simple paths from s to t. Here is a different but equivalent formulation of max-flow using flow paths. It has a variable x_p for every path p.

$$\boxed{\begin{array}{lll} \text{maximize} & \sum_{p\in P} f_P & \\ \text{subject to} & \displaystyle\sum_{p\in P: e\in P} f_p \le c_e & \forall e \in E \\ & f_p \ge 0 & \forall p \in P \end{array}} \tag{3}$$

Lemma: This LP is equivalent to LP (1) in the sense that they always have the same value.

Exer: Prove this lemma, by basically converting every feasible solution to one LP into a solution for the other LP.

The dual to LP (3) given below has variables $(y_e : e \in E)$.

$$\boxed{\begin{array}{lll} \text{minimize} & \displaystyle\sum_{e\in E} c_e y_e & \\ \text{subject to} & \displaystyle\sum_{e\in P} y_e \ge 1 & \forall p \in P \\ & y_e \ge 0 & \forall e \in E \end{array}} \tag{4}$$

This LP is an alternative relaxation of minimum cut.

Lemma: For every feasible solution $\vec{y}$ to LP (4) there is $A \subset V$, i.e. a cut $(A, \bar{A})$, such that

$$\text{capacity}(A, \bar{A}) \le \sum_{e\in E} c_e y_e.$$

Proof: seen in class, by interpreting y_e as edge lengths, computing shortest-paths distances from s and cutting at a random radius around s.

This gives yet another proof that

max-flow = fractional-cut = min-cut

KEY point: The dual of a flow problem involves distances, and distances are just "fractional" versions of cuts.

Exer: Turn the above proof into an algorithm that converts a solution to LP (4) into an st-cut with same value. What is its runtime?

Exer: Use a similar argument for the first relaxation of min-cut, i.e. prove that every solution to LP (2) can be converted into an st-cut with no larger value.

Exer: Does this proof work for undirected graphs?

2 Multicommodity flow and multicut

Let G be an undirected graph with edge-capacities, and for each $i \in [k]$ let $\{s_i, t_i\}$ be a pair of vertices called source-sink or demand pairs. The maximum multicommodity flow asks to ship the maximum amount of total flow, while respecting the edge capacities (the total flow on every edge is at most the edge's capacity).

The problem can be written as the LP below. We let P_i be the set of all $s_i - t_i$ paths.

$$\begin{array}{lll} \text{maximize} & \sum_i \sum_{p \in P_i} f_p^i & \\ \text{subject to} & \displaystyle\sum_{i \in [k]} \sum_{p \in P_i : e \in p} f_p^i \leq c_e & \forall e \in E \\ & f_p^i \geq 0 & \forall i \in [k], \forall p \in P_i \end{array} \tag{5}$$

In the multicut problem, the input is as above, and the goal is to find a minimum-capacity set of edges $E' \subset E$ such that removing E' from the graph disconnects every pair $s_i - t_i$.

Exer: Give an example where $G \setminus E'$ might have several connected components (for an optimal E').

Exer: show that in every network

$$\text{maximum multicommodity-flow} \leq \text{minimum multicut},$$

and give an example where the inequality is strict.

The dual of LP (5), given below, has variables $(y_e : e \in E)$.

$$\begin{array}{lll} \text{minimize} & \displaystyle\sum_{e \in E} c_e y_e & \\ \text{subject to} & \displaystyle\sum_{e \in p} y_e \geq 1 & \forall i \in [k], \forall p \in P_i \\ & y_e \geq 0 & \forall e \in E \end{array} \tag{6}$$

Observe that this LP is a relaxation of the multicut problem.

The flow/cut gap Theorem:

$$\text{minimum multicut} \leq O(\log k) \cdot \text{maximum multicommodity-flow}.$$

The proof idea is to again interpret the y_e variables as edge lengths, compute the shortest-path distance $d(u,v)$ between every two vertices, and "cut" around each t_i (or s_i). It relies on the following randomized for partitioning the vertices. For a partition P of V, we shall call each $S \in P$ a cluster.

Algorithm CKR/FRT
Input: metric d on V, a set $T \subset V$ of k vertices (called terminals), and parameter $R > 0$.
Output: a partition P of V
1. Choose a uniformly random permutation π of the terminals T
2. Choose $\beta \in (1,2)$ uniformly at random.
3. Iteratively for $t = \pi(1), \ldots, \pi(k)$, create a new cluster S_t containing the yet unclustered vertices u with $d(t,u) \le \beta R$.
4. Let P consist of these S_t plus all the remaining vertices as an additional cluster V_0.

Notice that t need not belong to "its own" cluster S_t.

Lemma: Every cluster $S_t \in P$ other than V_0 has diameter $< 4R$.

Proof: $d(u,v) \le d(u,t) + d(t,v) \le 2\beta R$.

Lemma: For every $u, v \in V$,

$$\Pr[u, v \text{ are separated in } P] \le O(\log k) \cdot d(u,v)/R.$$

Proof: seen in class, by analyzing the events of the form: a center t "captures" u but not v.

Remark (important for later): A more careful examination of the proof shows that

$$\Pr[u, v \text{ are separated in } P] \le \tfrac{d(u,v)}{R} O(\log \tfrac{|B_T(u,3R)|}{|B_T(u,R/2)|}).$$

The theorem can now be proved easily using the CKR algorithm.

Remark: It's not hard to verify that this gives a polynomial-time $O(\log k)$ approximation for the multicut problem, which is NP-hard even with $k = 3$.

Exer: Show a constant factor approximation for the multiway cut problem, which is a special case of multicut where there is a set T of k vertices (called terminals), and the demand pairs $\{s_i, t_i\}$ are just all $\binom{k}{2}$ pairs in T.