15-859(J): Metric Embeddings and Methods **Lecturer:** Anupam Gupta
Topic: Lower Bounds on Embeddings: Flows and Cuts **Date:** September 18, 2003
Scribe: Peter Richter

6.1 Generalized Flux and Multicommodity Flow

We have already seen that $(X, d) \xhookrightarrow{O(\log n)} \ell_1^{O(\log^2 n)}$. Let us relate the ℓ_1 embedding problem to the generalized flux problem to obtain a corresponding lower bound for distortion.

An instance of the *generalized flux problem* is a graph $G = (V, E)$ with *edge capacities* given by $\text{cap} : E \to \Re^+$ and pairwise *demands* given by $\text{dem} : V \times V \to \Re^+$. We will denote such an instance by $(G, \text{cap}, \text{dem})$. In Lecture 3, we defined the *generalized flux* across an edge cut ∂S in G be

$$\alpha_S = \frac{\text{cap}(\partial S)}{\text{dem}(S, V \setminus S)},$$

and the *generalized flux* α of the graph to be

$$\alpha = \min_{S \subseteq V} \alpha_S = \min_{\text{elementary cut metrics } d} \frac{\sum_e \text{cap}_e \, d_e}{\sum_{i,j} \text{dem}_{ij} \, d_{ij}} \tag{6.1.1}$$

Computing α using the right-hand-side identity above amounts to solving an integer program (IP). Let $S^* \subseteq V$ achieve the optimal flux, which we denote by $IP^* = \alpha_{S^*} = \alpha$. We can also relax the IP to obtain the following cut (primal) LP, by requiring that d merely be a metric, and not necessarily an elementary cut metric:

$$\begin{aligned} \min \quad & \sum_e \text{cap}_e \, d_e \\ \text{s.t.} \quad & \sum_{ij} \text{dem}_{ij} \, d_{ij} = 1 \\ & d_{ij} \le \sum_{e \in P} d_e \ \forall \text{ path } P \text{ between } i \text{ and } j, \forall i, j \\ & d_e \ge 0 \ \forall e \end{aligned} \tag{6.1.2}$$

Let the d^* be the metric that achieves the optimal value of this LP, and let the optimal value be denote by LP^*. The dual (or flow) LP is given by:

$$\begin{aligned} \max \quad & \lambda \\ \text{s.t.} \quad & \sum_{\text{P: P is a path between } i \text{ and } j} f_{ij}(P) \ge \lambda \, \text{dem}_{ij} \ \ \forall i, j \\ & \sum_{i,j} \sum_{P: e \in P} f_{ij}(P) \le \text{cap}_e \ \ \forall e \\ & \lambda \ge 0 \\ & f_{ij}(P) \ge 0 \ \forall \text{ path } P \text{ between } i \text{ and } j, \forall i, j \end{aligned} \tag{6.1.3}$$

Let the optimal value of the dual program be λ^*. Recall the following basic result from linear programming duality:

Theorem 6.1.1 (Strong duality) *The optimal values of the primal and dual linear programs are equal. I.e., $LP^* = \lambda^*$.*

On the other hand, $LP^* \leq IP^*$, since LP is a relaxation of the integer program IP. The inequality may be strict for some instances, such as the bipartite graph in Figure 6.1; here all the edge capacities (denoted by solid lines) and the symmetric demands (denoted by the dotted lines) are set to one.

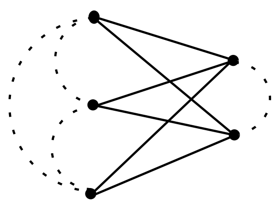

Figure 6.1.1: An instance in which $IP^* > LP^*$.

By separating any one of the 3 vertices from the larger set in the bipartition from the rest of the graph, we achieve $LP^* = \alpha = 1$. However, a simple volume argument gives us $\lambda^* \leq \frac{3}{4}$: each demand has to travel two edges to reach its destination, and hence every unit of demand sent must occupy two units of volume. Since each edge has unit capacity, and there are six edges, there are only six units of volume available. Hence $8\,\lambda \leq 6$, and thus $\lambda \leq \frac{3}{4}$. Hence, on this instance, $IP^* \geq \frac{4}{3}LP^*$.

Definition 6.1.2 *Given an integer program, and a linear programming relaxation for it, the* integrality gap *on an instance is the ratio IP^*/LP^* for that instance. The supremum of the ratio IP^*/LP^* over* all *instances is called the* integrality gap of the relaxation.

6.2 The Integrality Gap and Distortion into ℓ_1

The reason for our interest in the generalized flux problem is that its integrality gap is closely related to lower bounds for the distortion of embeddings into ℓ_1.

Theorem 6.2.1 *Given an instance $(G, \text{cap}, \text{dem})$ of the generalized flux problem, let d^* be the metric given by LP. If $(V, d^*) \overset{D}{\hookrightarrow} \ell_1$, then*

$$IP^* \leq D \cdot LP^*.$$

Proof: Theorem 6.2.1 The proof of this is a rephrasing of a proof given in Lecture 3, and is given here for completeness. We know that

$$LP^* = \sum_e \text{cap}_e\, d_e^* = \frac{\sum_e \text{cap}_e\, d_e^*}{\sum_{i,j} \text{dem}_{ij}\, d_{ij}^*}.$$

By our assumption, there is a metric μ in ℓ_1 such that

$$\mu \le d^* \le D \cdot \mu.$$

Hence,

$$LP^* \ge \frac{\sum_e \mathrm{cap}_e\, \mu_e}{D \cdot \sum_{i,j} \mathrm{dem}_{ij}\, \mu_{ij}}.$$

But since μ is an ℓ_1 metric, it is also a cut metric; i.e., $\mu_{ij} = \sum_{S \subseteq V : (i,j) \in \partial S} y_S$. So we have:

$$LP^* \ge \frac{\sum_S \mathrm{cap}(\partial S)}{D \cdot \sum_S y_S\, \mathrm{dem}(\partial S)} \ge \min_{S : y_S \ne 0} \frac{\mathrm{cap}(\partial S)}{D \cdot \mathrm{dem}(\partial S)} \ge \frac{IP^*}{D} \tag{6.2.4}$$

Therefore, $LP^* \le IP^* \le D \cdot LP^*$. ∎

Since we know that $D \le O(\log n)$ for any metric, we have the following corollary.

Corollary 6.2.2 *For all instances, $LP^* \le IP^* \le O(\log n) \cdot LP^*$.*

In other words, the best flow can be no better than the sparsest cut (worst bottleneck) and no worse than $O(\log n)$ times the sparsest cut.

Theorem 6.2.3 (Leighton-Rao (1988)) *For infinitely many values of n, there exist instances of the generalized flux problem on graphs G_n with n vertices such that*

$$IP^*/LP^* \ge \Omega(\log n).$$

Corollary 6.2.4 *Let the metric $d^*_{G_n}$ be generated by (LP) on the instances in Theorem 6.2.3. Then*

$$(V_n, d^*_{G_n}) \overset{\Omega(\log n)}{\hookrightarrow} \ell_1.$$

Proof: Suppose not, and let $(V_n, d^*_{G_n}) \overset{o(\log n)}{\hookrightarrow} \ell_1$. Then by Theorem 6.2.1, $IP^* \le o(\log n) \cdot LP^*$ on these instances, which contradicts Theorem 6.2.3. ∎

Hence, we have shown that there are metrics that require $\Omega(\log n)$ distortion to embed into ℓ_1, showing that Bourgain's embedding is existentially tight.

6.3 Maximizing the Integrality Gap

To wrap things up, we still need to prove Theorem 6.2.3. The lower bound is achieved by a family of expanders. Let us recall the following definition:

Definition 6.3.1 *The* expansion *(or sparsity) of set $S \subseteq V$ is given by:*

$$\Phi(S) := \frac{|\partial S|}{\min\{|S|, |V \setminus S|\}} \tag{6.3.5}$$

The expansion *$\Phi(G)$ of G is defined to be $\min_S \Phi(S)$. A graph G is an* expander *if the expansion is a constant independent of n; i.e., $\Phi(G) \ge \Omega(1)$.*

Constant-degree expanders are known to exist: this can be shown via a probabilistic argument. Explicit constructions are more difficult to come by, but these are also known:

Theorem 6.3.2 (Lubotzky, Phillips, and Sarnak (1988)) *For infinitely many n, there exist 3-regular expanders G_n.*

The expansion is closely related to the flux of the graph; note that if we set $\text{cap}_e = 1$ for all edges $e \in E$, and demands $\text{dem}_{ij} = 1$ for all pairs of vertices $i \neq j \in G$, we get:

$$\frac{n}{2} \cdot \alpha(G) \leq \Phi(G) \leq n \cdot \alpha(G) \tag{6.3.6}$$

Proof of Theorem 6.2.3: Let $G_n = (V_n, E_n)$ be a 3-regular expander, and fix any vertex $v \in V_n$. We claim that the number of vertices at distance at least $\frac{1}{2}\log_2 n$ from v is at least $\frac{1}{2}n$.

Indeed, in any Δ-regular graph, the number of vertices at distance no more than $t-1$ is at most $1 + \Delta + (\Delta-1)^2 + \ldots + (\Delta-1)^{t-1}$, which is

$$1 + \frac{(\Delta-1)^t - 1}{\Delta - 1}.$$

Now plugging in $\Delta = 3$ and $t = \frac{1}{2}\log n$, we get that the number of vertices *farther* than $\frac{1}{2}\log_2 n$ from v is at least $n - \sqrt{n} \geq n/2$.

Lemma 6.3.3 *For the 3-regular expanders of Theorem 6.3.2,*

$$LP^* \leq O(\frac{1}{n \log n}).$$

Proof: For each $v \in V_n$, look at all $\frac{1}{2}n$ or more vertices that are at distance at least $\frac{1}{2}\log_3 n$ from v. This gives at least $\frac{1}{2}n \cdot n$ demand pairs, such that sending λ units of flow between any one of these pairs uses up at least $\lambda \cdot \frac{\log_2 n}{2}$ units of volume. Since the total capacity available in the 3-regular graph is only $\frac{3}{2}n$, we must have that

$$\frac{1}{2}n \cdot n \cdot \frac{\lambda \log n}{2} \leq \frac{3}{2}n,$$

and hence $\lambda \leq \frac{6}{n \log n}$. ■

However, since G_n is an expander, $\Phi(G_n) = \Omega(1)$, and hence (6.3.6) implies that $\alpha(G_n) = IP^* \geq \Theta(\frac{1}{n})$. Therefore, the integrality gap $IP^*/LP^* \geq \Omega(\log n)$, which completes the proof of Theorem 6.2.3. ■

References

Leighton, F. T. and Rao, S. B. (1999). Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46(6):787–832. (Preliminary version in *29th FOCS*, pages 422–431, 1988).

Lubotzky, A., Phillips, R., and Sarnak, P. (1988). Ramanujan graphs. *Combinatorica*, 8(3):261–277.