On the optimal oblivious ratio of complete bipartite graphs

Stergios Athanassoglou * Jay Sethuraman †

April 2006

Abstract

The optimal oblivious ratio of complete bipartite graphs with unit edge capacities is shown to be bounded by 3.

1 Introduction

Traffic engineering is a critical component in managing large-scale networks. To manage such networks well, two related problems must be solved: estimating traffic flows, and designing effective routing protocols. If accurate estimation of traffic flows is possible, the routing problem is simply an instance of the well-studied multicommodity flow problem. However, obtaining an accurate estimate of traffic flows is often difficult, especially in a dynamic environment. Moreover, the routing protocol used may itself have an impact on the observed demand pattern, rendering the routing protocol ineffective. Recognizing these difficulties, recent research has focused on *oblivious routing*, which attempts to find a routing protocol that is robust with respect to traffic flows. An oblivious routing protocol employs a *fixed* routing regardless of the traffic pattern, and thus may perform worse than a cleverly-designed adaptive routing protocol; the advantage of an oblivious routing protocol is that it is significantly easier to implement.

Model. Let $G = (N, E)$ be a given *undirected* network, where N and E denote the node and edge sets of the network respectively. Let $c(e) > 0$ be the capacity of edge e. A routing f is a flow of one unit between all *ordered* pairs of nodes. The variable $f_{ij}(e)$ denotes the fraction of (the unit) flow from i to j that is routed on edge e. Let $D = [D_{ij}]$ be a demand matrix,

*IEOR Department, Columbia University, New York, NY; email: sa2164@columbia.edu
†IEOR Department, Columbia University, New York, NY; email: jay@ieor.columbia.edu

with D_{ij} representing the amount of traffic that must be sent from i to j. Denote the family of all possible demand matrices by $\mathcal{D}$. Given a routing f and a demand matrix $D \in \mathcal{D}$, the congestion, $EC(e, f, D)$, of edge e is simply

$$\sum_{i,j \in N, i \neq j} \frac{f_{ij}(e) D_{ij}}{c(e)}.$$

The network congestion, $C(f, D)$, is the maximum congestion incurred by any edge e. In other words,

$$C(f, D) = \max_{e \in E} EC(e, f, D).$$

Finally, given a demand matrix D, let $OPT(D)$ be the maximum congestion incurred by an offline optimal routing; as mentioned earlier, $OPT(D)$ can be found by solving a fractional multicommodity flow problem. Given these definitions, a reasonable measure of the performance of a routing f is the worst-case ratio of its congestion to that of the optimal congestion; this defines the oblivious performance ratio of f. The design problem then becomes one of finding a routing with the smallest oblivious performance ratio. Specifically, we wish to find a routing f^* such that

$$f^* = \arg\min_f \sup_{D \in \mathcal{D}} \left\{ \frac{C(f, D)}{OPT(D)} \right\}.$$

Related work and main result. In a remarkable paper Räcke [3] proved the existence of an oblivious routing for any undirected network that achieves a polylogarithmic competitive ratio with respect to congestion. Using a decomposition tree argument he was able to prove that one can always find a routing with an $O((\log n)^3)$ competitive ratio. Unfortunately, Räcke's algorithm relies on exponential-time procedures to build the decomposition tree, and so is not practical. This difficulty was corrected by Azar et al. [2], who propose a linear-programming based polynomial-time algorithm to find an optimal oblivious routing for any network. While this algorithm is polynomial-time, it relies on the ellipsoid method. Applegate and Cohen [1] propose a simpler linear-programming model using LP duality that avoids the ellipsoid method. They show that an optimal oblivious routing for a network with n nodes and m edges can be found by solving a single linear program with $O(mn^2)$ variables and $O(nm^2)$ constraints. Applegate and Cohen [1] consider cycles and complete graphs with unit-capacity edges. For these two classes of graphs, they show that the optimal oblivious ratio is $2 - 2/n$. Our main result is that for the complete bipartite graph $K_{m,n}$ with unit-capacity edges, the optimal oblivious ratio is

$$3 - \frac{2(m+n-1)}{mn};$$

moreover we construct a routing that achieves this performance ratio.

2 Main Result

Let $A = \{a_1, a_2, \ldots, a_m\}$ and $B = \{b_1, b_2, \ldots, b_n\}$ be the node sets of the bipartite graph. As in Applegate and Cohen [1], we assume the given graph is undirected, and there is a unit-capacity edge between a_i and b_j, for $i = 1, 2, \ldots, m$ and $j = 1, 2, \ldots n$. Without loss of generality, we assume $m \leq n$. The bounds shown here extend to the directed case as long as each (undirected) edge is replaced by a pair of anti-parallel directed edges, each with unit capacity. Since the graph is undirected, there exists a symmetric optimal oblivious routing [2, Lemma 8.1]: in such a routing $f_{ij}(e) = f_{ji}(e)$ for all pairs of nodes i and j, and all edges e. By this observation, we need only consider "undirected" commodities. In the rest of this note, the letters i and j are in the range $1, 2, \ldots, m$, and the letters k and l are in the range $1, 2, \ldots, n$, unless specified otherwise.

The rest of the section is devoted to a proof of the following theorem.

Theorem 1 *The optimal oblivious performance ratio for the complete bipartite graph $K_{m,n}$ with unit capacity edges is*

$$3 - \frac{2(m + n - 1)}{mn}.$$

Before turning to the proof, we make a few remarks that establish the *form* of an optimal routing. The possible commodities in our case naturally fall into one of two types, depending on whether or not the origin node lies on the same side as the destination node. The flow-conservation constraints and symmetry considerations imply that there is an optimal oblivious routing of the following form:

- For all $i \neq j$, $k \neq l$:

$$f_{a_i,b_k}(a_i, b_k) = x;$$

$$f_{a_i,b_k}(a_i, b_l) = \frac{1-x}{n-1}, \quad f_{a_i,b_k}(b_l, a_j) = \frac{1-x}{(m-1)(n-1)}, \quad f_{a_i,b_k}(a_j, b_k) = \frac{1-x}{m-1};$$

- For all $i \neq j$, for all k:

$$f_{a_i,a_j}(a_i, b_k) = f_{a_i,a_j}(b_k, a_j) = \frac{1}{n};$$

- For all i, for all $k \neq l$:

$$f_{b_k,b_l}(b_k, a_i) = f_{b_k,b_l}(a_i, b_l) = \frac{1}{m}.$$

Proof of Theorem 1.

Lower bound. Suppose there is a demand of one unit between a_i and b_k for all i, k. This traffic matrix is clearly routable with congestion 1 by sending each commodity along the direct edge. To find the congestion of routing f on this traffic matrix, we reason as follows. Any edge (a_i, b_k) carries four types of commodities:

(a) x units of the demand between a_i and b_k;

(b) $(1-x)/(n-1)$ units of the demand between a_i and b_l, for $l \neq k$;

(c) $(1-x)/(m-1)$ units of the demand between a_j and b_k, for $j \neq i$ and

(d) $(1-x)/\{(n-1)(m-1)\}$ units of the demand between a_j and b_l, for $j \neq i$, $l \neq k$.

Noting that there are $(n-1)$ commodities of type (b), $(m-1)$ commodities of type (c), and $(n-1)(m-1)$ commodities of type (d), we find that the edge (a_i, b_k) has a congestion of $3-2x$. In fact, this reasoning shows that the congestion on any given edge is $3-2x$ for this traffic matrix.

Now, consider the following traffic matrix: there is a demand of m units between a_1 and b_1; unit demand between the $(n-m)$ pairs $\{a_1, b_k\}$, $k = m+1, m+2, \ldots, n$; and unit demand between the $(m-1)(n-2)$ pairs $\{a_i, b_k\}$, for $i = 2, 3, \ldots, m$, $k = 2, 3, \ldots, n$, $k \neq i$. This traffic matrix is routable with congestion 1: send each of the unit demands on the direct edge; one unit of the demand between a_1 and b_1 along the direct edge; and the remaining $m-1$ units of demand between a_1 and b_1 along the $(m-1)$ paths (one unit on each path) (a_1, b_k, a_k, b_1) for $k = 2, 3, \ldots, m$. Under this traffic matrix, the routing f sends

$$mx + \frac{(1-x)}{n-1}(n-m) + \frac{1-x}{(m-1)(n-1)}(m-1)(n-2)$$

units of flow along the edge (a_1, b_1). Thus the congestion (largest flow on an edge) incurred in the graph is at least this quantity.

Since we do not know the traffic matrix, the worst-case congestion is at least

$$\max\left\{3-2x,\ mx + \frac{(1-x)}{n-1}(n-m) + \frac{1-x}{(m-1)(n-1)}(m-1)(n-2)\right\},$$

which is minimized when

$$3-2x \ = \ mx + \frac{(1-x)}{n-1}(n-m) + \frac{1-x}{(m-1)(n-1)}(m-1)(n-2).$$

Solving for x, we get:

$$x \ = \ \frac{m+n-1}{nm},$$

from which we infer a *lower bound* on the optimal oblivious ratio of at least $(3nm-2m-2n+2)/nm$.

Upper bound. Next we show that the routing f with $x = (m{+}n{-}1)/mn$ has an oblivious performance ratio that matches the lower-bound just described. This will simultaneously prove the tightness of the lower-bound and the optimality of f. For our specific choice of x, note that

$$\frac{1-x}{m-1} = \frac{n-1}{nm}, \quad \frac{1-x}{n-1} = \frac{m-1}{nm}, \quad \frac{1-x}{(m-1)(n-1)} = \frac{1}{nm}.$$

Consider any traffic matrix D that is routable with congestion 1. Fix an edge (a_i, b_k). We need to distinguish between six different kinds of commodities and their contribution to the congestion on edge (a_i, b_k): the commodity $\{a_i, b_k\}$; the $(n-1)$ commodities $\{a_i, b_l\}$ for $l \neq k$; the $(m-1)$ commodities $\{a_j, b_k\}$ for $j \neq i$; the $(n-1)(m-1)$ commodities $\{a_j, b_l\}$ for $j \neq i$, $l \neq k$; the $(m-1)$ commodities $\{a_i, a_j\}$ for $j \neq i$; and the $(n-1)$ commodities $\{b_k, b_l\}$ for $l \neq k$. It is easy to check that the total flow on edge (a_i, b_k) is:

$$xD_{a_i,b_k} + \frac{1-x}{n-1}\sum_{l:l\neq k} D_{a_i,b_l} + \frac{1-x}{m-1}\sum_{j:j\neq i} D_{a_j,b_k} + \frac{1-x}{(m-1)(n-1)}\sum_{j,l:j\neq i,l\neq k} D_{a_j,b_l} + \frac{1}{n}\sum_{j:j\neq i} D_{a_i,a_j} + \frac{1}{m}\sum_{l:l\neq k} D_{b_k,b_l},$$

which, using $x = (m+n-1)/mn$, simplifies to

$$\frac{(m+n-1)}{mn}D_{a_i,b_k} + \frac{m-1}{mn}\sum_{l:l\neq k} D_{a_i,b_l} + \frac{n-1}{nm}\sum_{j:j\neq i} D_{a_j,b_k} + \frac{1}{nm}\sum_{j,l:j\neq i,k\neq l} D_{a_j,b_l} + \frac{1}{n}\sum_{j:j\neq i} D_{a_i,a_j} + \frac{1}{m}\sum_{l:l\neq k} D_{b_k,b_l}. \quad (1)$$

To rewrite this expression in a more convenient way, it is useful to define

$$D(a_i) = D_{a_i,b_k} + \sum_{l:l\neq k} D_{a_i,b_l} + \sum_{j:j\neq i} D_{a_i,a_j},$$

and

$$D(b_k) = D_{a_i,b_k} + \sum_{j:j\neq i} D_{a_j,b_k} + \sum_{l:l\neq k} D_{b_k,b_l}.$$

Observe that $D(a_i)$ and $D(b_k)$ are the total demands involving nodes a_i and b_k respectively. For the given D matrix to be routable with congestion 1, it must be the case that

$$D(a_i) \leq n \quad \text{and} \quad D(b_k) \leq m,$$

because a_i has degree n and b_k has degree m. Using these definitions, the expression (1) for the congestion of the routing f can be rewritten as

$$\frac{m-1}{mn}D(a_i) + \frac{n-1}{mn}D(b_k) + \frac{1}{mn}\left\{D_{a_i,b_k} + \sum_{j,l:j\neq i,k\neq l} D_{a_j,b_l} + \sum_{j:j\neq i} D_{a_i,a_j} + \sum_{l:l\neq k} D_{b_k,b_l}\right\}.$$

To bound the last term, consider the cut given by $\{a_i, b_1, b_2, \ldots, b_{k-1}, b_{k+1}, \ldots, b_n\}$: its capacity is clearly $(m-1)(n-1)+1$, and each of the commodities that appear in the last term (within braces) has to "cross" this cut. We thus have shown

$$D_{a_i,b_k} + \sum_{j,l:j\neq i,k\neq l} D_{a_j,b_l} + \sum_{j:j\neq i} D_{a_i,a_j} + \sum_{l:l\neq k} D_{b_k,b_l} \;\leq\; (m-1)(n-1)+1.$$

Putting all this together, we infer that the total flow on edge (a_i, b_k) is at most

$$\frac{1}{mn}\left\{(m-1)(n) + (n-1)(m) + (m-1)(n-1) + 1\right\} \;=\; \frac{3mn-2m-2n+2}{mn} \;\leq\; 3.$$

■

References

[1] D. Applegate and E. Cohen, "Making Intra-Domain Routing Robust to Changing and Uncertain Traffic Demands: Understanding Fundamental Tradeoffs," *Proceedings of ACM SIGCOMM '03*, 313–324,2003.

[2] Y. Azar, E. Cohen, A. Fiat, H. Kaplan, and H. Racke, "Optimal oblivious routing in polynomial time," *Proceedings of the 35th ACM Symposium on the Theory of Computing*, June 2003.

[3] H. Räcke, "Minimizing congestion in general networks," *In Proceedings of the 43rd Annual Symposium on the Foundations of Computer Science*, pages 43–52, Nov. 2002