

Oblivious Routing on Node-Capacitated and Directed Graphs

MohammadTaghi Hajiaghayi

AT&T Labs—Research, Florham Park, New Jersey

Robert D. Kleinberg

Cornell University, Ithaca, New York

Harald Räcke

University of Warwick, Coventry, United Kingdom

and

Tom Leighton

Massachusetts Institute of Technology & Akamai Technologies, Cambridge, Massachusetts

Oblivious routing algorithms for general undirected networks were introduced by Räcke [Räcke 2002], and this work has led to many subsequent improvements and applications. Comparatively little is known about oblivious routing in general *directed* networks, or even in undirected networks with node capacities.

We present the first non-trivial upper bounds for both these cases, providing algorithms for k -commodity oblivious routing problems with competitive ratio $O(\sqrt{k} \log(n))$ for undirected node-capacitated graphs and $O(\sqrt{k} n^{1/4} \log(n))$ for directed graphs. In the special case that all commodities have a common source or sink, our upper bound becomes $O(\sqrt{n} \log(n))$ in both cases, matching the lower bound up to a factor of $\log(n)$. The lower bound (which first appeared in [Azar et al. 2003]) is obtained on a graph with very high degree. We show that in fact the degree of a graph is a crucial parameter for node-capacitated oblivious routing in undirected graphs, by providing an $O(\Delta \text{polylog}(n))$ -competitive oblivious routing scheme for graphs of degree Δ . For the directed case, however, we show that the lower bound of $\Omega(\sqrt{n})$ still holds in low-degree graphs.

Finally, we settle an open question about routing problems in which all commodities share a common source or sink. We show that even in this simplified scenario there are networks in which no oblivious routing algorithm can achieve a competitive ratio better than $\Omega(\log n)$.

Categories and Subject Descriptors: C.2.2 [Computer-Communication Networks]: Network Protocols—*Routing protocols*; F.2.0 [Analysis of Algorithms and Problem Complexity]: General

General Terms: Algorithms

Additional Key Words and Phrases: Oblivious routing; Node-Capacitated graphs; Directed graphs; Communication networks

A preliminary version of this paper appeared in the *Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms*, 2005.

The work was done while M.T. Hajiaghayi and R.D. Kleinberg were in the MIT Computer Science and Artificial Intelligence Laboratory and H. Räcke was in the Department of Computer Science, Carnegie Mellon University. R.D. Kleinberg was supported by a Fannie and John Hertz Foundation Fellowship.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2008 ACM 1529-3785/2008/0700-0111 \$5.00

1. INTRODUCTION

A routing algorithm for large-scale, unstructured networks like the Internet has to meet many partly conflicting criteria, e.g., enabling quick routing decisions, performing well under a variety of different traffic patterns, and working in a distributed fashion in order to keep control overhead low. In particular the latter issue creates serious difficulties, because a lack of coordination due to distributed routing decisions can easily create high-loaded “hot spots” within the network that usually results in a very bad routing performance.

In this paper we consider the online virtual circuit routing problem in a general network. In this problem a sequence of routing requests is received in an online manner. Each request consists of a pair of nodes that wish to communicate, and a routing algorithm has to establish for each request a path through the network that connects the source to the target. We consider two cost-measures, namely the edge-congestion and the node-congestion, which are defined as the maximum load of a network edge and network node, respectively. (The load of an edge (a node) is the number of traversing paths divided by the capacity of the edge (the node).) This cost-measure reflects the goal of minimizing “hot spots”, and thereby prevents the emergence of bottlenecks in the network.

Aspnes et al. [Aspnes et al. 1997] give an algorithm for edge-congestion that achieves an optimum competitive ratio of $O(\log n)$, where n denotes the number of vertices in the network. Since their algorithm works for directed graphs, a standard reduction from node-capacitated to directed graphs can be used to obtain the same competitive ratio for node congestion, as well. Unfortunately, this result does not properly address the need for distributed routing decisions, as the algorithm is *adaptive*, i.e., its decisions depend on the current load in the network. Therefore Aspnes et al. assume that there is a centralized decision maker that has full instantaneous knowledge of the traffic situation in the network. This can only be implemented with a lot of control overhead.

For the cost measure of edge-congestion in undirected graphs this problem has been recently solved by [Räcke 2002] at the cost of a slightly higher competitive ratio. It is shown that for any undirected network, there is a routing algorithm with polylogarithmic competitive ratio that is oblivious, i.e., the routing decisions are independent from the current load in the network and a *fixed* flow routing is used for any set of demands.

1.1 Our contribution

In this paper we analyze the performance of oblivious routing algorithms for the cost-measure of edge-congestion in directed graphs and for the cost-measure of node-congestion in undirected graphs.¹

As shown by Azar et al. [Azar et al. 2003], it is not in general possible to get a polylogarithmic competitive ratio in directed graphs due to a lower bound of $\Omega(\sqrt{n})$. We extend this lower bound to undirected node-capacitated graphs, i.e., to the case of

¹A simple reduction shows that in directed graphs the cost-measures node congestion and edge congestion are equivalent. Therefore, in the rest of the paper when we consider node capacities we implicitly assume that the underlying graph is undirected.

Type of oblivious routing	Lower bound	Upper bound
undirected edge-capacitated (single-sink)	$\Omega(\log n)$ [this paper]	$O(\log^2 n \log \log n)$ [Harrelson et al. 2003] (previously $O(\log^3 n)$ [Räcke 2002])
undirected edge-capacitated (general case)	$\Omega(\log n)$ [Bartal and Leonardi 1999]	$O(\log^2 n \log \log n)$ [Harrelson et al. 2003] (previously $O(\log^3 n)$ [Räcke 2002])
undirected node-capacitated (single-sink)	$\Omega(\sqrt{n})$ [this paper]	$O(\min(\sqrt{n} \log n, \Delta \log^2 n \log \log n))$ [this paper]
undirected node-capacitated (general case)	$\Omega(\sqrt{n})$ [this paper]	$O(\min(\sqrt{k} \log^{3/2} n, \Delta \log^2 n \log \log n))$ [this paper]
bounded degree directed (single-sink)	$\Omega(\sqrt{n})$ [this paper]	$O(\sqrt{n} \log n)$ [this paper]
directed (general case)	$\Omega(\sqrt{n})$ [Azar et al. 2003]	$O(\sqrt{kn}^{1/4} \log n)$ [this paper]
directed with symmetric demands	$\Omega(\sqrt{n})$ [this paper]	$O(\sqrt{k} \log^{5/2} n)$ [this paper]

Table I. Competitive ratios for different types of oblivious routing. Here Δ denotes the maximum degree of the graph and k denotes the potential number of commodities ($k = O(n^2)$ in the worst case)

node congestion in undirected graphs.

Then we show that for node capacitated graphs the critical parameter is not the number n of graph nodes but the maximum degree Δ of a node in the network by providing an algorithm that obtains competitive ratio $O(\Delta \text{polylog } n)$. This is done via a reduction to the undirected edge-capacitated case. We show that in general this result can not be improved by more than a polylogarithmic factor as there exist graphs with maximum degree $\Theta(\sqrt{n})$ in which every oblivious routing scheme has competitive ratio $\Omega(\sqrt{n})$. It is worth mentioning that such a result cannot be obtained for directed graphs, since we show that even for directed graphs of degree at most three, the lower bound on the competitive ratio is $\Omega(\sqrt{n})$.

Unfortunately, the algorithm may be far away from optimum if the maximum degree in the graph is $\Delta = \Theta(n)$. Furthermore, it seems to be very difficult to transfer the technique of using a hierarchical decomposition (and in fact constructing such a decomposition) to the case of undirected node-capacitated graphs to get a better result. This seems to be even harder for directed graphs.

Therefore we introduce a new approach to oblivious routing that gives a substantial improvement in the competitive ratio on high-degree node-capacitated networks and on

directed graphs. The performance of our algorithm depends on two parameters. The first one is k ; the maximum number of commodities that potentially may be routed. This means that the adversary chooses demands only from a restricted predetermined set of at most k source-sink pairs (where k is at most n^2). The second parameter is the maximum possible ratio between the throughput of a maximum concurrent flow and the capacity of a sparsest cut. We call this parameter the max-flow min-cut gap and denote it with α . Note that α depends on the underlying graph and on the set of source-sink pairs for which the adversary may create a demand. For example, if we consider general commodities in an undirected (edge-capacitated or node-capacitated) graph $\alpha = O(\log n)$ [Garg et al. 1996; 2004]; if we consider flow problems with symmetric demands in directed graphs (symmetry means that for any two nodes u, v the demand from u to v is equal to the demand from v to u) $\alpha = O(\log^3 n)$ [Even et al. 1998; Klein et al. 1997]. For these cases our algorithm obtains a competitive ratio of $O(\sqrt{\alpha k} \log n) = O(\sqrt{k} \text{polylog } n)$.

Another important implication of our oblivious routing algorithm in directed graphs is that we obtain competitive ratio $O(\sqrt{n} \log n)$ when there is a common source or a common sink that is shared by all commodities.² The single-source case arises naturally, for instance, in communication networks where all clients are receiving their data from a single server such as an HTTP or a streaming-media server. Our result in this case is tight up to a logarithmic factor due to the lower bound of $\Omega(\sqrt{n})$ by Azar et al. [Azar et al. 2003].

As a further result, we disprove the existence of a constant upper bound for single-sink oblivious routing in undirected edge-capacitated graphs by providing an $\Omega(\log n)$ lower bound (the upper bound for this case is $O(\log^2 n \log \log n)$, which is the upper bound for the general multicommodity case [Harrelson et al. 2003]). This answers a question asked by Azar et al. [Azar et al. 2003] for oblivious routing with a common source or sink.

The reader is referred to Table 1 to see a complete comparison between the results of this paper and the previous work.

1.2 Related work

The idea of selecting routing paths oblivious to the traffic in the network has been intensively studied for special network topologies, since such algorithms allow for very efficient implementations due to their simple structure. Valiant and Brebner [Valiant and Brebner 1981] initiate the worst case theoretical analysis for oblivious routing on the hypercube. They design a randomized packet routing algorithm that routes any permutation in $O(\log n)$ steps. This result gives a virtual circuit routing algorithm that obtains a competitive ratio of $O(\log n)$ with respect to edge-congestion.

In [Räcke 2002] it is shown that there is an oblivious routing algorithm with polylogarithmic competitive ratio (w.r.t. edge-congestion) for any undirected graph. However, this result is non-constructive in the sense that only an exponential time algorithm was given for constructing the routing scheme.

²In the remainder of the paper we will usually refer to this scenario as the *single-sink case*. However, all results for a single sink hold for a single source, as well.

This issue was subsequently addressed by Azar et al. [Azar et al. 2003] who show that the optimum oblivious routing scheme, i.e., the scheme that guarantees the best possible competitive ratio, can be constructed in polynomial time by using a linear program. This result holds for edge-congestion, node-congestion and in arbitrary directed and undirected graphs. Furthermore, they show that there are directed graphs such that every oblivious routing algorithm has a competitive ratio of $\Omega(\sqrt{n})$.

The method used by Azar et al. does not give the possibility to derive general bounds on the competitive ratio for certain types of graphs. Another disadvantage of [Azar et al. 2003] is that it does not give a polynomial time construction of the hierarchy used in [Räcke 2002], which has proven to be useful in many applications (see e.g. [Alon et al. 2004; Chekuri et al. 2004; Maggs et al. 2005]). A polynomial time algorithm for this problem was independently given by [Bienkowski et al. 2003] and [Harrelson et al. 2003]. Whereas the first result shows a slightly weaker competitive ratio for the constructed hierarchy than the non-constructive result in the original paper, the second paper by Harrelson, Hildrum and Rao has even improved the competitive ratio to $O(\log^2 n \log \log n)$. This is currently the best known bound for oblivious routing in general undirected graphs.

Other distributed routing and admission control algorithms are also proposed by Awerbuch and Azar [Awerbuch and Azar 1994] and Awerbuch and Leighton [Awerbuch and Leighton 1993; 1994] which route flows with a rate that is within a $(1 + \epsilon)$ factor of the optimal, but these are not real-time algorithms and take at least a polylogarithmic number of rounds to converge.

2. FORMAL DEFINITION OF THE PROBLEM

We represent the network as a graph $G = (V, E)$ (directed or undirected), where V denotes the set of vertices (or nodes) and E denotes the set of edges. We denote the number of vertices by n . The degree of a vertex v is denoted by $d(v)$, and the maximum degree of a node in G by Δ . We will assume that a capacity function cap is given, assigning a capacity (or bandwidth) to either nodes or edges in the graph. This models the physical communication potential of the network resources.

For this network, we are further given a set K of commodities (let $k = |K|$). Each commodity $(s, t) \in K$ specifies a source node $s \in V$ and a target node $t \in V$ that potentially want to communicate. An oblivious routing scheme for K specifies a unit flow between source and target for each commodity in K . The unit flow for a commodity $(s, t) \in K$ defines a “routing rule” that describes how demand between source s and target t is routed through the network. In the rest of the paper, when we do not specify K explicitly, we implicitly assume that K is the set of all n^2 node pairs.

For a given set of demands and a given routing algorithm, we define the *absolute load* of an edge (a node) as the amount of data routed along this edge (the node). (Note that a message that starts at node $v \in V$ will contribute to the load of v .) The relative load is the absolute load of an edge or a node divided by the respective capacity. The *node-congestion* is the maximum relative load of any node in the network, and the *edge-congestion* is defined as the maximum relative load of an edge.

The goal is to design an oblivious routing algorithm that always achieves an edge or node congestion that is close to the best possible. To formalize this we introduce

the notion of a demand vector D that specifies the demand for every commodity in K . We denote the optimum edge congestion and optimum node congestion that can be obtained for demands D with $\text{OPT}_E(D)$ and $\text{OPT}_V(D)$, respectively. When there is no ambiguity we use $\text{OPT}_E(D)$ and $\text{OPT}_V(D)$ also to denote the actual routing that achieves the optimum congestion.

For an oblivious routing strategy OBL, let $\text{congestion}_E(\text{OBL}, D)$ and $\text{congestion}_V(\text{OBL}, D)$ denote the edge congestion and node congestion, respectively, achieved for demand vector D . The competitive ratio of the oblivious routing scheme is defined as

$$c_E(\text{OBL}) := \max_D \left\{ \frac{\text{congestion}_E(\text{OBL}, D)}{\text{OPT}_E(D)} \right\}$$

and

$$c_V(\text{OBL}) := \max_D \left\{ \frac{\text{congestion}_V(\text{OBL}, D)}{\text{OPT}_V(D)} \right\},$$

for edge congestion and node congestion, respectively.

3. OBLIVIOUS ROUTING ON DIRECTED GRAPHS

We consider oblivious routing algorithms for a directed graph $G = (V, E)$ with a set K of k commodities given in advance. The capacity of an edge $e \in E$ is denoted by $\text{cap}(e)$. For a commodity $i \in K$, the source and the target for commodity i will be denoted by s_i and t_i respectively. The demand for commodity i will be denoted by d_i . We will assume that the commodities and their demand patterns belong to a class for which the max-flow min-cut gap is at most α , where α is a parameter known to the algorithm and depends on the specifics of the type of the flow problem being considered. Thus, for instance, $\alpha = 1$ if we are considering single-sink or single-source flow problems, $\alpha = O(\log n)$ [Garg et al. 1996; 2004] if we are considering general multicommodity flows in undirected graphs, $\alpha = O(\log^3 n)$ [Even et al. 1998; Klein et al. 1997] if we are considering multicommodity flow problems with symmetric demands in directed graphs, and $\alpha = O(\sqrt{n})$ [Hajiaghayi and Räcke] if we are considering general multicommodity flows in directed graphs.

THEOREM 3.1. *There is an oblivious routing algorithm OBL that achieves competitive ratio $c_E(\text{OBL}) = O(\sqrt{\alpha k} \log n)$ on directed graphs.*

PROOF. The high level idea for the proof is as follows. We partition the set of commodities into subsets K_1, K_2, \dots and bound the load created on an edge by any of these subsets individually. (Each subset may use a different routing strategy). We show that if the demands of a particular subset can be routed (by an optimal algorithm) with congestion C , i.e., $C = \text{OPT}_E(D)$, an edge will get load at most $\sqrt{\alpha k} \cdot C$ from such a subset. Moreover, each edge gets non-negligible load from only $O(\log n)$ of the subsets. This means that the oblivious routing algorithm is $O(\sqrt{\alpha k} \log n)$ -competitive.

A subset contains commodities that have approximately the same mincut. Formally let $K_j, j = 0, 1, \dots$ denote the set of all commodities i for which the capacity of the minimum cut in G separating s_i from t_i is in the interval $[2^j, 2^{j+1})$. (Note that we can assume without loss of generality that the minimum capacity of an edge is 1.)

We now describe how commodities in a specific set K_j are routed. For this we split the set K_j into two subsets A_j and B_j as described by the following lemma.

LEMMA 3.2. *A set K_j of commodities can be split into two possibly empty subsets A_j and B_j such that*

- (1) *the commodities in A_j can be well routed, i.e., each source can send a flow of 2^j to its target concurrently with total congestion $\leq \sqrt{\alpha \cdot k}$, and*
- (2) *the commodities in B_j are separated by a small cut, i.e., there is a set S_j of edges that separate every commodity in B_j and has $\text{cap}(S_j) \leq \sqrt{\alpha/k} \cdot |B_j| \cdot 2^j$.*

These sets may be computed in polynomial time, given an oracle which takes a flow problem as input and produces a cut whose sparsity is within a factor α of the maximum concurrent flow.

PROOF. We construct sets A_j, B_j and the cut S_j as follows. We initialize B_j and S_j to be empty and A_j to be equal to K_j . Consider the flow problem where each commodity in A_j sends a flow of 2^j to its target. If this flow problem has congestion larger than $\sqrt{\alpha k}$, the α -approximate max-flow min-cut theorem ensures that there is an edge-cut S' that separates a subset $A' \subset A_j$ of commodities such that $|A'|2^j/\text{cap}(S') > \sqrt{\alpha k}/\alpha = \sqrt{k/\alpha}$. Now we have

$$\begin{aligned} \text{cap}(S_j \cup S') &\leq \text{cap}(S_j) + \text{cap}(S') \\ &\leq \sqrt{\alpha/k} \cdot |B_j| \cdot 2^j + \sqrt{\alpha/k} \cdot |A'| \cdot 2^j \\ &= \sqrt{\alpha/k} \cdot |B_j \cup A'| \cdot 2^j, \end{aligned}$$

hence moving the set A' from A_j to B_j and adding the edges in S' to S_j generates B_j and S_j that fulfill the second requirement of the theorem. Iterating this procedure generates a set A_j that fulfills the first requirement. Hence the lemma follows. \square

Now we specify how commodities in sets A_j and B_j are routed. The commodities in A_j are essentially routed according to the solution of the concurrent multicommodity flow problem with demand 2^j for each commodity (by construction we know that the congestion of this flow problem is only $\sqrt{\alpha k}$). However, we modify the flow solution, without increasing the congestion of any edge by a factor of more than 2, so that it does not use any edge of capacity less than $2^j/2n^4$. (Let E_{low} denote the set of edges with capacity less than $2^j/2n^4$. The total demand that is shipped over these edges is at most $|E_{\text{low}}| \cdot \frac{2^j}{2n^4} \cdot \sqrt{\alpha k} \leq n^2 \cdot \frac{2^j}{2n^4} \cdot n^2 \leq 2^j/2$. Each commodity $i \in A_j$ sends demand 2^j ; hence less than half of its flow is routed along flow paths containing an edge of capacity less than $2^j/2n^4$. Reroute all such flow, by at most doubling the flow on each path which avoids such low-capacity edges.) This defines the routing scheme for commodities in A_j . For a commodity $i \in B_j$, the routing is determined by computing a congestion-minimizing single commodity flow for i , then modifying it as above so that the congestion of each edge increases by a factor of at most 2, and no edges of capacity $< 2^j/2n^2$ are utilized.

Note that we can scale the above flows for commodities in A_j and B_j to value 1 so that we obtain an oblivious routing scheme.

We claim that this oblivious routing scheme for the commodity set K_j is $O(\sqrt{\alpha k})$ -competitive for flow problems limited to that commodity set. If $\{d_i : i \in K_j\}$ is a demand pattern admitting a flow solution with congestion C , then $d_i \leq C2^{j+1}$ for each i , so the commodities in A_j place congestion at most $4C\sqrt{\alpha k}$ on any edge. Turning our attention now to the commodities in B_j , let $W = \sum_{i \in B_j} d_i$ denote the total demand for commodities in B_j . Each commodity $i \in B_j$ places a congestion of at most $2d_i/2^j$ on any edge. (This is because i is routed according to a 2-approximation to the congestion-minimizing single-commodity flow, and the mincut for a commodity is at least 2^j which means that there is a single-commodity flow for i with congestion at most $d_i/2^j$.) Consequently, the total congestion on an edge due to commodities in B_j is at most $2W/2^j$. But $W \leq C\text{cap}(S_j)$, since each commodity in B_j is separated by the cut S_j and there exists a flow solution with congestion C . This gives, for each edge e ,

$$\begin{aligned} \text{congestion}(e) &\leq \frac{2W}{2^j} \\ &\leq \frac{2C\text{cap}(S_j)}{2^j} \\ &\leq 2C\sqrt{\alpha/k} \cdot |B_j| \\ &\leq 2C\sqrt{\alpha k} . \end{aligned}$$

We have now seen, for each edge e , that if the optimal flow solution has congestion C , then in the oblivious flow solution each commodity set K_j will place congestion at most $6C\sqrt{\alpha k}$ on edge e . It remains to bound the number of commodity sets which place a non-negligible load on e . A set K_j only uses an edge e if $\text{cap}(e) \geq 2^j/2n^4$. Let j_{\max} denote the largest j such that this holds, and let $j' = j_{\max} - 6\log_2 n - 1$. If $j' < 0$ then e only gets load from $O(\log n)$ sets K_j . Otherwise we have $\text{cap}(e) \geq 2^{j_{\max}}/2n^4 = 2^{j'}n^2 \geq 2^{j'}k$. This means if every commodity in sets K_j , $j < j'$ sends its entire demand (at most $C \cdot 2^{j'}$) along e , the resulting congestion on e only increases by C . This means the total congestion on e is only $O(C + (j_{\max} - j')\sqrt{\alpha k}C) = O(C \cdot \sqrt{\alpha k} \log n)$, as desired.

4. OBLIVIOUS ROUTING ON UNDIRECTED NODE CAPACITATED GRAPHS

In this section, we present an upper bound $O(\Delta \text{polylog } n)$ for the performance ratio of oblivious routing on undirected node-capacitated graphs. We note that this upper bound can only be improved by a logarithmic factor, since in the construction mentioned in the proof of Theorem 5.2 (see Section 5), the lower bound is $\Omega(\sqrt{n}) = \Omega(\Delta)$.

THEOREM 4.1. *There is an oblivious routing algorithm OBL that achieves competitive ratio $c_V(\text{OBL}) = O(\Delta \log^2 n \log \log n)$ on undirected node-capacitated graphs with maximum degree Δ .*

PROOF. The proof follows from a reduction from the node-capacitated case to the edge-capacitated case and then uses the result of Räcke [Räcke 2002] (and its improvement by Harrelson, Hildrum, and Rao [Harrelson et al. 2003]) for the edge-capacitated case.

From the given undirected graph $G = (V, E)$, we construct a new graph $G' = (V', E')$ as follows. For a node $v \in V$ with degree $d(v)$, we create $d(v)$ vertices $v_1, v_2, \dots, v_{d(v)}$ in V' (we say vertices $v_1, v_2, \dots, v_{d(v)}$ in V' are clones of vertex v in V). Edges are defined as follows. We order the neighbors of a vertex v from 1 to $d(v)$ arbitrarily, and place an edge between v_i and u_j if and only if v is the i -th neighbor of u and u is the j -th neighbor of v in the aforementioned orderings. The capacity of such an edge is ∞ . In addition, we connect the vertices $v_1, v_2, \dots, v_{d(v)}$ in V' by a clique. The capacity of each edge in this clique is $c(v)/d(v)$.

Now, we construct the oblivious routing on G' using the algorithm of Harrelson et al. [Harrelson et al. 2003] with competitive ratio $O(\log^2 n \log \log n)$ and obtain the final oblivious routing in G , by contracting all vertices $v_1, v_2, \dots, v_{d(v)}$ into one vertex v (and thus a path which goes through clones of a vertex $v \in V$ in G' goes through v in G .) Next, we show that the competitive ratio of the resulting oblivious routing is good.

Consider a demand vector (matrix) D and the corresponding optimal routing $\text{OPT}_V(D)$ for node congestion. We show that we can route the corresponding demand matrix D' in G' with edge congestion at most $2\text{OPT}_V(D)$ (here by corresponding demand matrix we mean for any demand d from u to v in G , there is a demand d from an arbitrary copy u_i , $1 \leq i \leq d(u)$, to an arbitrary copy v_j , $1 \leq j \leq d(v)$). Consider a path $P = (u = v^1, v^2, \dots, v^k = v)$ in G which carries a non-zero fraction of the demand from u to v ; say this fraction is ϵ . We consider the corresponding path P' in G' between the clones of vertices v^1, v^2, \dots, v^k . We only have to analyze the load on clique-edges as other edges have infinite capacity. Fix three consecutive nodes v^{i-1}, v^i, v^{i+1} in the path and assume that the path P' enters the clique corresponding to v^i at node v_j^i and leaves it at v_k^i (i.e., v^{i-1} is the j -th neighbor of v^i and v^{i+1} is the k -th neighbor). To forward the ϵ -flow from clone v_j^i to v_k^i we split it among all paths $v_j^i - x - v_k^i$ where x is an intermediate node of the clique that is different from v_j^i (note that x may be equal to v_k^i in which case we use the direct link between the nodes).

Hence, for any ϵ -fraction of demand that goes through $v \in V$ a clique edge only gets an $\epsilon/(d(v)-1)$ -fraction. Since the capacity of each edge is $c(v)/d(v)$, it means the total congestion of edges in the clique is at most $\frac{d(v)\beta}{d(v)-1} \leq 2\beta$, where β is the congestion of v in $\text{OPT}_V(D)$. Since the edges between clones of different nodes $u, v \in V$ have infinite capacity, the total congestion of the solution in G' corresponding to $\text{OPT}_V(D)$ in G is at most $2 \cdot \text{OPT}_V(D)$.

Now if we use the oblivious routing of Harrelson et al. [Harrelson et al. 2003] in G' the resulting edge congestion is at most $O(\log^2 n \log \log n)$ of the congestion of the solution constructed from $\text{OPT}_V(D)$ and thus at most $2 \cdot O(\log^2 n \log \log n) \cdot \text{OPT}_V(D)$. Now, when we shift from the oblivious routing in G' to the oblivious routing in G , we have a blow-up factor of at most $(\Delta - 1)/2$. This is because when we contract all clones of a vertex in order to obtain the final oblivious routing in G , the load of a vertex v is bounded above by the load on all its clique-edges which is at most $\frac{d(v)(d(v)-1)}{2} \cdot \frac{c(v)}{d(v)} \cdot \gamma$, where γ is the maximum congestion of an edge in the clique. Hence the congestion of the node is at most $O((d(v)-1)\gamma)$. This means the overall congestion of a node using our oblivious routing in G is at most $O(\Delta \log^2 n \log \log n) \cdot \text{OPT}_V(D)$ as desired. \square

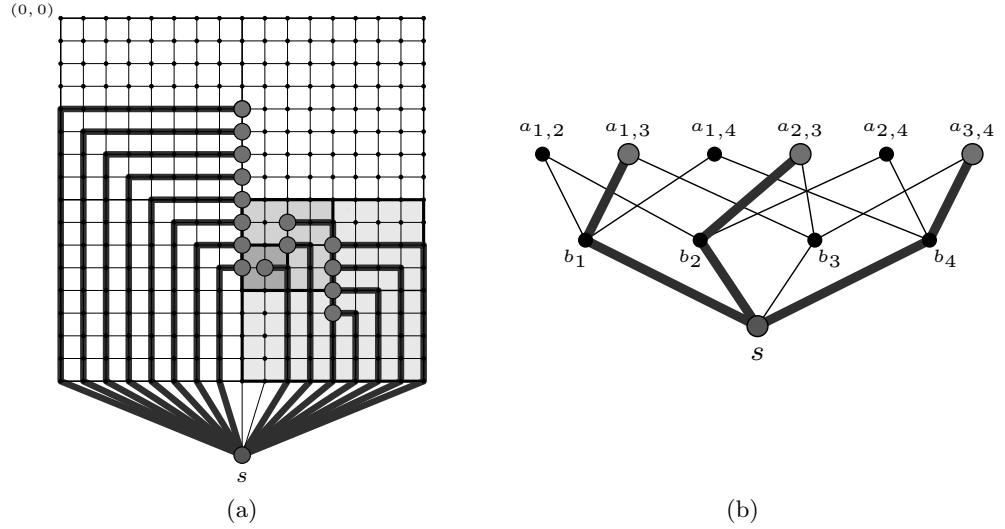


Figure 1. (a) The mesh M with a common sink used in Theorem 5.1. The gray nodes represent a possible set of demand-nodes according to the recursive construction. The routing paths give an optimal routing for these demand nodes. (b) The network for proving the lower bound for node-capacitated oblivious routing.

5. LOWER BOUNDS

In this section, we first disprove the existence of a constant upper bound for single-sink oblivious routing in undirected edge-capacitated graphs by providing an $\Omega(\log n)$ lower bound.

THEOREM 5.1. *There is an undirected graph G such that for any oblivious routing algorithm OBL, $c_E(\text{OBL}) = \Omega(\log n)$, even if all commodity pairs share a common sink.*

PROOF. The proof is an extension of the proof by Bartal and Leonardi [Bartal and Leonardi 1999] and Maggs et al. [Maggs et al. 1997] for general multicommodity online routing. Consider an $N \times N$ 2-dimensional mesh M , with $N = 2^k + 1$ rows and columns for some integer $k \geq 1$. We number rows and columns from 0 to 2^k and use $M[x, y]$ to denote the vertex in row x and column y . We add a super-sink s which is connected to all vertices in the bottom row, i.e., vertices $M[2^k, y]$, $0 \leq y \leq 2^k$. We consider the scenario in which each node of M is a possible source that may communicate with s . All edge-capacities are one. Figure 1a shows M for $k = 4$.

The request sequence is defined recursively in k stages as follows. At the i th stage of the recursion where $i = 0, 1, \dots, k - 1$, we define the requests for a $(2^{k-i} + 1) \times (2^{k-i} + 1)$ sub-mesh M_i determined from the previous step. We set $M_0 = M$ in the beginning. In the i th stage each vertex $M_i[x, 2^{k-i-1}]$ where $2^{k-i-2} \leq x < 3 \cdot 2^{k-i-2}$ issues a request of demand one. The number of edges of the mesh M_i is $2 \cdot 2^{k-i} \cdot (2^{k-i} + 1)$. At the i th stage, the path for a demand-node uses at least 2^{k-i-2} edges of sub-mesh M_i , because

the distance from the demand-node to the boundary of M_i is at least that large. The number of requests at this stage is at least 2^{k-i-1} . Thus the total load induced by these requests on sub-mesh M_i is at least $2^{k-i-2} \cdot 2^{k-i-1}$. Hence the average load on edges of M_i is at least $\frac{2^{k-i-2} \cdot 2^{k-i-1}}{2 \cdot 2^{k-i} \cdot (2^{k-i} + 1)} \geq 1/32$.

We partition M_i into four submeshes of size $(2^{k-i-1} + 1) \times (2^{k-i-1} + 1)$ such that each submesh contains a corner of M_i . An edge is contained in at most two of these sub-meshes. Therefore, one of the submeshes has average load at least $((1/32)/2)/4 = 1/256$. This submesh is chosen as submesh M_{i+1} in the following stage. We can do this recursively; in each step we take the sub-mesh that currently has highest average load according to the routing of the oblivious scheme and we recursively generate new demand within this sub-mesh. By induction we can prove that the average load of edges in the sub-mesh M_i is at least $i/256$. This means there exists an edge in M_{k-1} with load $\Omega(\log n)$, where $n = N^2 + 1$ denotes the number of nodes in the graph.

On the other hand, an optimal algorithm can in each step use those edges in M_i which are not in M_{i+1} to route the demands of M_i to the sink. In this way, all edges of M_{i+1} are free of load and we can use them for the requests in the subsequent stages (Figure 1a shows an example for $k = 4$).

Thus, we can choose the requests in such a way that the oblivious routing algorithm has congestion $\Omega(\log n)$, while an optimum solution routes all requests with congestion one. This finishes the proof. \square

Note that the above lower bound does not only hold for oblivious routing schemes but for any online algorithm (in an online model the routing requests arrive sequentially and an online algorithm can base its routing decisions on past requests; thus an online algorithm is more powerful than an oblivious routing scheme). In addition, Theorem 5.1 shows that the result of Harrelson et al. [Harrelson et al. 2003] who give a competitive ratio of $O(\log n \log \log n)$ for oblivious routing on planar graphs is tight up to a factor of $O(\log \log n)$ even for single-sink oblivious routing (since the graph considered above is planar).

The following theorem gives a lower bound of $\Omega(\sqrt{n})$ for the single-sink oblivious routing problem in general undirected node-capacitated graphs. This shows that our upper bound of $O(\sqrt{n} \log n)$ given in Section 3 is tight up to a logarithmic factor.

THEOREM 5.2. *There is an undirected graph G such that for any oblivious routing algorithm OBL, $c_V(\text{OBL}) = \Omega(\sqrt{n})$, where the set of commodity pairs shares a common sink.*

PROOF. We essentially use the same example introduced by Azar et al. [Azar et al. 2003]. It consists of a graph G with three levels. The first level contains $\binom{k}{2}$ nodes denoted by a_{ij} for $1 \leq i < j \leq k$. The second level contains k nodes denoted by b_i for $1 \leq i \leq k$ and the third level contains a super-sink s . Each node a_{ij} is connected via two undirected edges to nodes b_i and b_j . Furthermore, each node b_i is connected to the sink. All nodes have capacity one, except the sink s which has infinite capacity.

Any oblivious routing scheme defines a unit flow from each node a_{ij} on the first level to the sink. Hence, there is a node b_x which receives at least $\binom{k}{2}/k$ units of flow from its direct neighbors on the first level (i.e., nodes a_{ix} for $i < x$ and nodes a_{xj} for $j > x$).

Suppose that all these neighbors of b_x send a demand of one and all other demands are zero. Then the oblivious algorithm has congestion at least $\binom{k}{2}/k$ at node b_x . However, the optimum algorithm can route this demand with congestion 1 by using the paths $a_{ix} - b_i - s$ for demands from nodes a_{ix} and the paths $a_{xj} - b_j - s$ for demands from a_{xj} -nodes. The proof follows immediately, since $k = \Omega(\sqrt{n})$. Figure 1b illustrates the case for $k = 4$ and $x = 3$. \square

The following theorem shows that the result of Theorem 4.1, which provides an $O(\Delta \text{polylog } n)$ -competitive oblivious routing algorithm for node-capacitated undirected graphs, cannot be extended to directed graphs.

THEOREM 5.3. *There is a directed graph G of degree at most three such that for any oblivious routing algorithm OBL, $c_E(\text{OBL}) = \Omega(\sqrt{n})$, where the set of commodity pairs share a common sink.*

PROOF. Consider the directed version of the graph used in the proof of Theorem 5.2, i.e., each edge is directed towards the higher level node and is assigned a unit capacity. Now, we replace each node b_i and the sink s by a directed tree, in which edges are directed towards the root and have infinite capacity. The graph edges that are directed to a node b_i or to the sink s are attached to distinct leaf nodes of the corresponding tree, and the outgoing edge of node b_i is attached to the root. Since all flow that is sent to such a b_i -tree has to traverse the outgoing edge attached to the root, the analysis of Theorem 5.2 works for the new graph. Since the degree of each node in the new graph is at most three, the theorem follows. \square

Finally, we note that the construction considered in the proof of Theorem 5.2 also shows that the performance of oblivious routing with symmetric demands in edge-capacitated directed graphs is at least $\Omega(\sqrt{n})$. To see this, consider the directed version of the graph G in which all edges are directed towards the higher level node, and add edges with infinite capacity from s to all nodes on the first level. In this graph symmetry of demands does not change the problem and therefore the $\Omega(\sqrt{n})$ lower bound holds.

6. DISCUSSION AND OPEN PROBLEMS

In this paper, we presented non-trivial upper and lower bounds for oblivious routing in node-capacitated and directed graphs. In particular, we obtained almost tight upper and lower bounds (up to an $O(\log n)$ factor) for single-sink oblivious routing.

The main open problem is whether we can obtain a competitive ratio of $O(\sqrt{n} \text{polylog } n)$ for oblivious routing in general node-capacitated graphs when the number k of potential commodities is asymptotically larger than n (in this case, we have an upper bound of $O(\sqrt{k} \text{polylog } n)$). The problem for directed graphs seems more challenging, since the best known max-flow min-cut gap is $O(\sqrt{n})$ ([Hajiaghayi and Räcke]) in this case (In all currently known algorithms for oblivious routing this max-flow min-cut gap plays an important role).

The results of Harrelson et al. [Harrelson et al. 2003] for the competitive ratio of oblivious routing in edge-capacitated undirected graphs are parameterized in terms of the max-flow min-cut gap α of the considered graph-class. They obtain an

$O(\alpha \log n \log \log n)$ -competitive oblivious routing algorithm. On planar graphs, e.g., α is constant. Unfortunately, our reduction from the node-capacitated case to the edge-capacitated case does not preserve planarity. Hence, it is an interesting open problem whether for planar graphs an $O(\Delta \log n \log \log n)$ -competitive oblivious routing algorithm can be obtained for the node-capacitated version of the problem.

Last but not least, there is a slight gap of $O(\log n)$ between the lower bounds and the upper bounds for the single-sink case in both directed and undirected graphs. Closing these gaps is another interesting open problem.

REFERENCES

- ALON, N., AWERBUCH, B., AZAR, Y., BUCHBINDER, N., AND NAOR, J. S. 2004. A general approach to online network optimization problems. In *Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 577–586.
- ASPINES, J., AZAR, Y., FIAT, A., PLOTKIN, S. A., AND WAARTS, O. 1997. On-line routing of virtual circuits with applications to load balancing and machine scheduling. *Journal of the ACM* 44, 3, 486–504. Also in *Proc. 25th STOC*, 1993, pp. 623–631.
- AWERBUCH, B. AND AZAR, Y. 1994. Local optimization of global objectives: competitive distributed deadlock resolution and resource allocation. In *Proceedings of the 35th IEEE Symposium on Foundations of Computer Science (FOCS)*. 240–249.
- AWERBUCH, B. AND LEIGHTON, F. T. 1993. A simple local-control approximation algorithm for multicommodity flow. In *Proceedings of the 34th IEEE Symposium on Foundations of Computer Science (FOCS)*. 459–468.
- AWERBUCH, B. AND LEIGHTON, F. T. 1994. Improved approximation algorithms for the multicommodity flow problem and local competitive routing in dynamic networks. In *Proceedings of the 26th ACM Symposium on Theory of Computing (STOC)*. 487–496.
- AZAR, Y., COHEN, E., FIAT, A., KAPLAN, H., AND RÄCKE, H. 2003. Optimal oblivious routing in polynomial time. In *Proceedings of the 35th ACM Symposium on Theory of Computing (STOC)*. 383–388.
- BARTAL, Y. AND LEONARDI, S. 1999. On-line routing in all-optical networks. *Theoretical Computer Science* 221, 1-2, 19–39. Also in *Proc. 24th ICALP*, 1997, pp. 516–526.
- BIENKOWSKI, M., KORZENIOWSKI, M., AND RÄCKE, H. 2003. A practical algorithm for constructing oblivious routing schemes. In *Proceedings of the 15th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. 24–33.
- CHEKURI, C., KHANNA, S., AND SHEPHERD, B. 2004. The all-or-nothing multicommodity flow problem. In *Proceedings of the 36th ACM Symposium on Theory of Computing (STOC)*. 156–165.
- EVEN, G., NAOR, J. S., SCHIEBER, B., AND SUDAN, M. 1998. Approximating minimum feedback sets and multicuts in directed graphs. *Algorithmica* 20, 2, 151–174.
- GARG, N., VAZIRANI, V. V., AND YANNAKAKIS, M. 1996. Approximate max-flow min-(multi)cut theorems and their applications. *SIAM Journal on Computing* 25, 2, 235–251.
- GARG, N., VAZIRANI, V. V., AND YANNAKAKIS, M. 2004. Multiway cuts in node weighted graphs. *Journal of Algorithms* 50, 1, 49–61.
- HAJIAGHAYI, M. T. AND RÄCKE, H. An $O(\sqrt{n})$ -approximation algorithm for directed sparsest cut. *Information Processing Letters*. To appear.
- HARRELSON, C., HILDRUM, K., AND RAO, S. B. 2003. A polynomial-time tree decomposition to minimize congestion. In *Proceedings of the 15th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. 34–43.
- KLEIN, P. N., PLOTKIN, S. A., RAO, S. B., AND TARDOS, É. 1997. Approximation algorithms for Steiner and directed multicuts. *Journal of Algorithms* 22, 2, 241–269.

- MAGGS, B. M., MEYER AUF DER HEIDE, F., VÖCKING, B., AND WESTERMANN, M. 1997. Exploiting locality for networks of limited bandwidth. In *Proceedings of the 38th IEEE Symposium on Foundations of Computer Science (FOCS)*. 284–293.
- MAGGS, B. M., MILLER, G. L., PAREKH, O., RAVI, R., AND WOO, S. L. M. 2005. Finding effective support-tree preconditioners. In *Proceedings of the 17th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*. 176–185.
- RÄCKE, H. 2002. Minimizing congestion in general networks. In *Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science (FOCS)*. 43–52.
- VALIANT, L. G. AND BREBNER, G. J. 1981. Universal schemes for parallel communication. In *Proceedings of the 13th ACM Symposium on Theory of Computing (STOC)*. 263–277.