Semi-oblivious Routing: Lower Bounds

MohammadTaghi Hajiaghayi* Robert Kleinberg† Tom Leighton‡

Abstract

We initiate the study of *semi-oblivious* routing, a relaxation of oblivious routing which is first introduced by Räcke and led to many subsequent improvements and applications. In semi-oblivious routing like oblivious routing, the algorithm should select only a polynomial number of paths between the source and the sink of each commodity, but unlike oblivious routing, the flow from each source to its sink is not just a scalar multiple of the single-commodity flow; any amount of flow can be sent along each selected path. Semi-oblivious routing has several applications in traffic engineering and VLSI routing.

Trivially, any competitive ratio ρ for oblivious routing (includling the polylogarthimic ratio in undirected graphs obtained by Räcke) also implies competitive ratio ρ for semi-oblivious routing. In this paper, we focus on lower bounds. We rule out the possibility of $O(1)$ competitive ratio for semi-oblivious routing in undirected graphs by providing a lower bound of $\Omega(\frac{\log n}{\log \log n})$ in grids or even series-parallel graphs. More strongly in directed graphs, we rule out the possibility of sub-polynomial competitive ratio when the number of paths between each source and its sink is in $O(n^{1/5})$. The proof of our lower bound on the grid uses a non-Markovian random walk on the integers with a mixing property which may be of independent interest. Last but not least, our lower bounds on the grid can be significantly strengthened to show that with paths of at most b bends, the competitive ratio is in $\Omega(n^{\frac{1}{2b+1}})$. This answers negatively a long-standing open problem on b-bend routing schemes in grids posed e.g. in [10, 6].

1 Introduction

In this work we initiate the study of *semi-oblivious* routing problems, a class of flow problems which can be described by the following series of interactions between an algorithm and an adversary. First, the algorithm selects a set Π of paths in a specified graph G of size n. The size of Π must be polynomial in n. Second, the adversary presents a multicommodity flow problem in G. Third, the algorithm computes a solution to the specified multicommodity flow problem. Every unit of flow must be routed using one of the paths in Π. The objective is to minimize the congestion of the most congested edge or vertex in G. We refer to the set of paths Π as a *semi-oblivious routing scheme* in G. We say that Π is *ρ-competitive* if it is the case that for every demand matrix that might be selected in step 2, the algorithm can compute a flow in step 3 whose congestion is a ρ-approximation to the optimum congestion.

Semi-oblivious routing problems are analogous to oblivious routing problems [11], which may be described using the same pattern of steps as above, except that in step 1 the algorithm chooses a single-commodity flow for each pair s, t (instead of choosing a set of $s-t$ paths) and the flow routed from s to t in step 3 must use a scalar multiple of the single-commodity flow chosen in step 1. It is known [5, 8, 11] that every undirected graph has an oblivious routing scheme which is $O(\log^2 n \log\log n)$-competitive with respect to edge congestion, and that such a routing scheme can be computed in polynomial time. However, it is also known [4, 7] that in general graphs — in fact, even in grids — the optimal oblivious routing scheme is $\Omega(\log n)$-competitive, and that in directed graphs or node-capacitated undirected graphs the competitive ratio of the optimal oblivious routing scheme can be as large as $\Omega(\sqrt{n})$. Can some of these lower bounds be overcome by using semi-oblivious routing rather than oblivious routing? In particular, can we achieve competitive ratio $O(1)$ in undirected graphs, or at least in grids? We settle this question negatively in Theorems 3.1 and 3.3 below.

Semi-oblivious routing is interesting not only as a relaxation of oblivious routing, but also as an abstraction of several problems arising in network traffic engineering and in VLSI design. We briefly touch on each of these applications next.

Traffic engineering. A common practice in modern Internet traffic engineering is the use of multiprotocol label switching (MPLS) to create explicit label-switched paths (LSPs) between certain ingress and egress points in an autonomous system. By properly apportioning traffic among these *LSP tunnels*, ISP's can achieve desired resource utilization and performance objectives [2]. Because the instantiation and maintenance of LSP tunnels consumes network

*Computer Science Department, Carnegie Mellon University. Supported in part by NSF ITR grant CCR-0122581 (The Aladdin Center). hajiagha@cs.cmu.edu

†Computer Science Division, UC Berkeley; and Computer Science Department, Cornell University. Supported by an NSF Mathematical Sciences Postdoctoral Research Fellowship. rdk@cs.cornell.edu

‡Department of Mathematics, Massahusetts Institute of Technology; and Akamai Technologies. ftl@math.mit.edu

management resources, it is infeasible to maintain a large (e.g. exponential in the network size) set of LSPs. Our semi-oblivious routing problem thus abstracts the problem one faces when trying to instantiate LSP tunnels to allow unknown future demand matrices to be routed in a way which will approximately minimize congestion.

VLSI design. Routing in the grid to minimize congestion has been extensively studied, largely due to its applications in the layout of VLSI circuits. Routing with a constant number of bends was studied in [1, 9]; from these works, it is known that for routing problems in which each node of the $n \times n$ grid belongs to at most one source-sink pair, the congestion of the optimal one-bend routing is bounded above by $\lceil n/2 \rceil + 1$, and that a routing with $b+1$ bends can improve significantly (i.e. by a factor of $n^{\Omega(1)}$) on the congestion of the optimal routing with b bends when $b = 1, 2$. Note that for constant b, the set of all b-bend paths in the grid is a semi-oblivious routing scheme with $O(n^{b+3})$ paths, so these results can all be interpreted as results about semi-oblivious routing in the grid. Our Theorem 3.1 thus places a non-trivial lower bound on the congestion that can be achieved using routes with $b = O(1)$ bends in the $n \times n$ grid. For b-bend routing schemes, this lower bound can be significantly improved; see Theorem 3.4 below. This answers negatively an old open problem on b-bend routing schemes on grids posed e.g. in [10, 6].

2 Definitions

Throughout this paper, "log" denotes the base-2 logarithm.

Definition Let G be a graph with n vertices and m edges, Π a set of paths in G, and let $\Pi(s,t)$ denote the set of all paths in Π which start at s and end at t. A Π-valued flow in G is a vector $\mathbf{f}$ specifying a non-negative weight f_P for each $P \in \Pi$. We say that $\mathbf{f}$ *satisfies* a non-negative demand matrix $D = (D_{st})_{s,t \in V(G)}$ if it is the case that

$$D_{st} = \sum_{P \in \Pi(s,t)} f_P$$

for all $s, t \in V(G)$.

Suppose that each edge e of G has a positive real-valued capacity $c(e)$. We say that $\mathbf{f}$ has *edge congestion* C if

$$C = \max_{e \in E(G)} \frac{1}{c(e)} \sum_{P \in \Pi \,:\, e \in P} f_P.$$

Similarly, if each vertex v of G has a positive real-valued capacity $c(v)$, we say that $\mathbf{f}$ has *vertex congestion* C if

$$C = \max_{v \in V(G)} \frac{1}{c(v)} \sum_{P \in \Pi \,:\, v \in P} f_P.$$

Definition A *semi-oblivious routing scheme* in a graph G is a set Π of paths in G, such that for each pair of distinct vertices s, t, Π contains at least one path from s to t. Suppose that each edge e of G has a positive real-valued capacity $c(e)$. Given a demand matrix D, we define $\mathrm{opt}_\Pi(D)$ to be the minimum edge congestion of a Π-valued flow which satisfies D. We define the *competitive ratio* of Π to be

$$\rho(\Pi) = \sup_D \left(\frac{\mathrm{opt}_\Pi(D)}{\mathrm{opt}_{\mathbb{P}(G)}(D)} \right),$$

where $\mathbb{P}(G)$ denotes the set of all paths in G and the supremum is taken over all nonzero demand matrices D.

3 Our results

Trivially, if there exists an oblivious routing scheme in a graph G with competitive ratio ρ, then there exists a semi-oblivious routing scheme with competitive ratio ρ. (Simply take each of the flows in the oblivious routing scheme, and represent it as a weighted sum of paths.) Thus, using the results of [5, 8, 11], every undirected graph of size n has a semi-oblivious routing scheme whose competitive ratio (with respect to edge-congestion) is $O(\log^2 n \log\log n)$. Our first result shows that this upper bound can not be improved to $O(1)$, even when the graph is a grid.

THEOREM 3.1. *The competitive ratio of every semi-oblivious routing scheme in the $n \times n$ grid is in $\Omega(\log n / \log\log n)$.*

The proof of Theorem 3.1 uses a non-Markovian random walk on the integers with a mixing property which may be of independent interest.

THEOREM 3.2. *There exists a random walk $\ldots, Z_{-1}, Z_0, Z_1, \ldots$ on the integers with the following property. For every integer $\ell > 1$, every residue class $x \in \mathbb{Z}/(\ell)$, and every pair of integers t, u such that $t - u > 8\ell \log^3(2\ell)$, $\Pr(Z_t \equiv x \pmod{\ell} \parallel Z_u) > \frac{1}{2\ell}$.*

Given the lower bound for semi-oblivious routing in grids, perhaps there is a simpler class of graphs (e.g. graphs of bounded treewidth) which admits semi-oblivious routing schemes with congestion $O(1)$. The following lower bound for series-parallel graphs (which have treewidth 2) rules out this possibility.

THEOREM 3.3. *There exist series-parallel graphs of size n in which every semi-oblivious routing scheme has competitive ratio $\Omega(\log n / \log\log n)$.*

The lower bound in Theorem 3.1 can be strengthened when the semi-oblivious routing scheme in question consists of all paths with at most b bends, for some constant b.

THEOREM 3.4. *There exists a multicommodity flow instance in the n-by-n grid which can be satisfied by a flow with edge-congestion one, but any flow which satisfies the* *demands using paths with less than b bends has congestion $\Omega(n^{\frac{1}{2b+1}})$.*

Theorem 3.4 is especially interesting since answers negatively an old open problem (see e.g. [10, 6]) which asks whether or not allowing only a constant number of bends per each flow path in multicommodity routing on the grid is sufficient to achieve at most a constant factor (even a polylogarithmic factor) times optimal congestion.

Finally the main question left open by our work is the following.

Question Does every directed graph G with n vertices have a semi-oblivious routing scheme (with a polynomial number of paths) with congestion $O(\text{polylog}(n))$? Does every node-capacitated undirected graph have such a routing scheme?

The following theorem provides evidence for a negative answer to this question, though the number of paths between each commodity is still sub-linear.

THEOREM 3.5. *There is a node-capacitated undirected graph with n vertices such that if Π is any semi-oblivious routing scheme satisfying $|\Pi(s,t)| = O(n^{1/5})$ for all s,t, then the competitive ratio of Π, with respect to vertex congestion, is $\Omega(n^{1/10})$. A fortiori, the same lower bound holds for directed graphs.*

Theorems 3.1 and 3.2 are proved in Section 5, Theorem 3.3 in Section 4, Theorem 3.4 in Section 6, and Theorem 3.5 in Section 7.

4 Semi-oblivious routing in series-parallel graphs

In this section we describe the construction which underlies Theorem 3.3, showing that series-parallel graphs do not have semi-oblivious routing schemes which use a polynomial number of paths and achieve competitive ratio $o(\log n/\log\log n)$. Our construction uses a family of graphs $G(k,r)$ defined recursively as follows. Let $K = k^6$ and let $[K] = \{1, 2, \ldots, K\}$. $G(k,1)$ consists of a sequence of vertices $s = s_0, s_1, \ldots, s_K = t$ with a pair of parallel edges joining s_{i-1} to s_i for $1 \le i \le K$. (If a simple graph is desired, one can subdivide each of the edges into a path of length two. The lower bound will remain valid.) Recursively define $G(k,r)$ by starting with a copy of $G(k,r-1)$ and transforming each edge $e = (u,v)$ into a copy of $G(k,1)$, identifying the vertices s,t of $G(k,1)$ with the endpoints u,v of edge e. The graphs $G(k,1)$ and $G(k,2)$ are illustrated in Figure 1.

A random multicommodity flow in $G(k,r)$, which we will denote by $F(k,r)$, can be defined recursively as follows. In $G(k,1)$ the flow $F(k,1)$ chooses a random path from s to t and routes one unit of flow along this path. There is a complementary path from s to t whose edges carry no flow. In the recursive step, we are given a random multicommodity flow $F(k,r-1)$ in $G(k,r-1)$ such that there is a path from s to t whose edges carry no flow, and all other edges carry one unit of flow. We first double all the demands in $F(k,r-1)$. When we build $G(k,r)$ from $G(k,r-1)$ we replace each edge $e = (u,v)$ of $G(k,r-1)$ with a copy of $G(k,1)$. If edge e carried a unit of flow in $F(k,r-1)$ then we route two units of flow from u to v by sending one unit on each edge of the copy of $G(k,1)$. The remaining edges of $G(k,r-1)$ become a subgraph of $G(k,r)$ consisting of K copies of $G(k,1)$ concatenated serially. In each of these copies of $G(k,1)$ we route one unit of flow from the vertex which corresponds to s to the vertex which corresponds to t, using the random flow $F(k,1)$. Let $D_{k,r}$ be the random demand matrix satisfied by the flow $F(k,r)$.

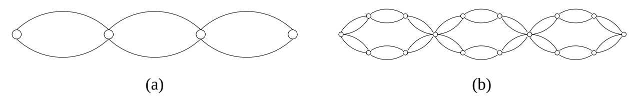

Figure 1: The graphs $G(3,1)$ and $G(3,2)$.

Setting $r = k$, we have a defined a random flow problem in $G(k,k)$. The lower bound of $\Omega(k) = \Omega(\log n/\log\log n)$ comes about from the following observations; those which are not obvious from the construction are proved below.

1. There are $k = \Omega(\log n/\log\log n)$ different classes of commodities, distinguished by the distance between their source and sink.

2. The optimal solution routes all the commodities in such a way that each edge is used by only one class of commodities, and the congestion on that edge is 1. In fact, the partition of the edge set among the commodity classes is achieved as follows: first half of the edges are chosen at random and reserved for class 1. Then, recursively, the remaining edges are partitioned among the remaining classes.

3. Let L be the total load induced on all edges by the optimum flow. The combined capacity of all edges in G is $O(L)$. The load from routing commodities in class j (for $1 \le j \le k$) is $2^{-j}L$.

4. A small fraction (i.e., 2^{-k}) of the edges are not used by any class of commodities in the optimal routing solution. Call this edge set E_*. Unfortunately, for each j these edges are located so close to the sources and sinks for commodities of type j that with high probability, all of the paths used by the routing scheme Π for those commodities will have about 2^{j-k} fraction of their edges in E_*. Consequently the amount of load on edges in E_* caused by commodities in class j is about $2^{-k}L$.

5. Summing over all commodity classes, the edges in E_* carry a load of $k \cdot 2^{-k}L$. Since their capacity is only $O(2^{-k}L)$, there must be an edge in E_* whose congestion is k.

Before we begin the proof, let us specify the counterexample more precisely. As above, given an integer $k > 1$, let $K = k^6$. The edge set E of $G(k,r)$ is the set $[K^r] \times \{0,1\}^r$ (edges of the graph are arranged in K^r levels from left to right, as shown in Figure 1.) The vertex set V of $G(k,r)$ is a quotient set of $\{0,1,\dots,K^r\} \times \{0,1\}^r$. Define an equivalence relation $\sim$ on this set by specifying that for two elements (p,a) and (q,b), the relation $(p,a) \sim (q,b)$ holds if and only if $p = q = jK^\ell$ for some non-negative integers j,ℓ, and $a_i = b_i$ for $\ell+1 < i \le r$. The vertex set of G is the set of equivalence classes of $\sim$, and the left and right endpoints of an edge whose label is (p,a) are the equivalence classes containing $(p-1,a)$ and (p,a), respectively. The equivalence class of $(0,0)$ will be denoted by s, and the equivalence classes of $(K^r,0)$ will be denoted by t.

A path P from s to t in $G(k,r)$ is specified by a sequence of edges $(m,a^{(m)})$, where m runs from 1 to K^r. The bit strings $a^{(m)}$ must satisfy $a_i^{(m-1)} = a_i^{(m)}$ for all $i > 1$ such that $K^{i-1} \not| m$. For any path P from s to t, let $\psi_i(P) \in \{0,1\}^{K^r}$ denote the sequence $(a_i^{(1)}, a_i^{(2)}, \dots, a_i^{(K^r)})$. (Here i can be any number between 1 and r.) As we have seen, the sequence $\psi_i(P)$ is constant between consecutive powers of K^{i-1}.

For any path P from s to t, specified by a sequence of edges $(m,a^{(m)})$ as above, a partition of the edge set $E \setminus P$ into subsets E_i $(1 \le i \le r)$ is defined as follows. Edge $e = (m,a)$ belongs to E_i if $a_i \ne a_i^{(m)}$ but $a_j = a_j^{(m)}$ for all $j < i$. Clearly $|E_i| = |E|/2^i$. A demand matrix D^P is defined as follows. $D^P_{v,w} = 2^{r-i}$ if there exist integers $p,q \in \{0,\dots,K^r\}$ such that p is a multiple of K^{r+1-i}, $q = p + K^{r+1-i}$, v is the equivalence class of $(p-1,a^{(p)})$, and w is the equivalence class of $(q,a^{(q)})$. Otherwise $D^P_{v,w} = 0$. When P is a path from s to t chosen uniformly at random, the distribution of the demand matrix D^P is the same distribution denoted by $D_{k,r}$ in Section 4. For a commodity with source $v = (p-1,a)$ and sink $w = (q,b)$, we define the *class* of the commodity to be the number $r + 1 - \log_K(q-p)$. Thus, for example, class 1 consists solely of the commodity with source s and sink t.

If v,w is a commodity in class i, it is easy to see that E_i contains 2^{r-i} disjoint paths of length K^{r+1-i} joining v to w. The demand $D^P_{v,w}$ may be routed from v to w by sending one unit of flow on each of these paths. In this way, we can define a flow satisfying the entire demand matrix D^P, such that the flow for commodities in class i put one unit of load on all the edges in E_i and does not use any other edges. We call this flow the *optimal flow*. It puts one unit of load on $L = (1-2^{-r})|E|$ of the edges of G.

Let $F_i = E_{i-1} \setminus E_i$. (When $i = 1$ we are adopting the convention that $E_0 = E$. We will also put $F_0 = E$.) We have $P = F_r \subset F_{r-1} \subset \ldots \subset F_1 \subset F_0 = E$, and each edge set has twice as many elements as the preceding one.

LEMMA 4.1. *If Q is a path joining the source and sink of a commodity in class i, then at least K^{r+1-i} edges of Q belong to F_{i-1}.*

Proof. [Proof sketch.] When $i = 1$ there is nothing to prove, since $F_0 = E$. When $i > 1$, the reader may verify from the definition of F_{i-1} that it contains every shortest path from v to w. But the structure of $G(k,r)$ is such that any path from v to w which is not a shortest path is a very circuitous path that has even more than K^{r+1-i} edges in F_{i-1}. See Figure 1 for intuition. □

LEMMA 4.2. *Suppose P is a uniformly random path in $G(k,r)$. Let us condition on the sequences $\psi_1(P),\dots,\psi_{i-1}(P)$; this conditioning will be implicit in all expressions involving probabilities and expectations. (Note that the conditioning determines the edge set F_{i-1}.) Let Q be any simple path in $G(k,r)$ joining two vertices v,w such that $v = (p-1,a)$, $w = (q,b)$, and both p and q are divisible by K^{r+1-i}. Then $\mathbb{E}(|Q \cap F_i|) = \frac{1}{2}|Q \cap F_{i-1}|$, and*

$$\Pr\left(|Q \cap F_i| < \left(\frac{1}{2} - \frac{1}{k^2}\right)|Q \cap F_{i-1}|\right) < e^{-2k^2}.$$

Proof. If $Q \cap F_{i-1} = \emptyset$ then there is nothing to prove. Otherwise, the structure of $G(k,r)$ guarantees that any simple path which intersects F_{i-1} does so in a union of subpaths of length K^{r+1-i}. Each of these subpaths is further partitioned into K subpaths of equal length, each of which is either disjoint from F_i or contained in it according to the value of one bit of the sequence $\psi_i(P)$. Since different subpaths correspond to different bits, we find that the random variable $(|Q \cap F_i|)/(|Q \cap F_{i-1}|)$ has the same distribution as the fraction of heads observed in at least K tosses of a fair coin. The expression for $\mathbb{E}(|Q \cap F_i|)$ is now obvious, and the tail bound for $|Q \cap F_i|$ follows by applying Chernoff's bound to the fraction of heads observed in at least K fair coin tosses. □

LEMMA 4.3. *If Π is any set of $e^{o(k^2)}$ paths in $G(k,k)$, and P is a uniformly random path in $G(k,r)$, then with probability $1-o(1)$ the following holds for every path $Q \in \Pi$ joining the source and sink of a commodity in class j of the flow problem D^P:*

$$|Q \cap P| \ge (1-1/k)2^{j-k}|Q|.$$

Proof. Lemma 4.1 says that $|Q \cap F_{j-1}| \ge |Q|$. Applying Lemma 4.3 for $i = j, j+1, \dots, k$, we find that

$$\Pr\left(|Q \cap P| < \left(\frac{1}{2} - \frac{1}{k^2}\right)^{j-k}|Q|\right) < ke^{-2k^2}.$$

Using the fact that $\left(\frac{1}{2} - \frac{1}{k^2}\right)^{j-k} > \left(1 - \frac{1}{k}\right) 2^{j-k}$, we obtain $\Pr\left(|Q \cap P| < \left(1 - \frac{1}{k}\right)\right) < ke^{-2k^2}$. The lemma follows by applying the union bound over all paths in Π. □

THEOREM 4.1. *Any semi-oblivious routing scheme* Π *in* $G(k,k)$ *has competitive ratio* $\Omega(k) = \Omega(\log n / \log\log n)$.

Proof. We have $|\Pi| = n^{O(1)} = e^{O(k \log k)}$, so we may apply Lemma 4.3 to conclude that with high probability over the random choice of a demand matrix D^P, the inequality $|Q \cap P| \geq (1 - 1/k)2^{j-k}|Q|$ holds for all paths $Q \in \Pi$ joining the source and sink of a commodity with positive demand in D^P. Assume from now on that this holds for all such Q. The commodities in class j $(1 \leq j \leq k)$ contribute $\Omega\left(2^{-j}|E|\right)$ units of load in the optimal flow by sending all of their demand along shortest paths. Therefore, in any Π-valued flow satisfying D^P, the commodities in class j also contribute $\Omega\left(2^{-j}|E|\right)$ units of load, and at least $\Omega\left(2^{j-k}2^{-j}|E|\right) = \Omega\left(2^{-k}|E|\right)$ of this load is on edges of P. Summing over commodity classes $1, 2, \ldots, k$, we see that the total load on edges of P is $\Omega\left(k2^{-k}|E|\right)$. Since $|P| = 2^{-k}|E|$, it follows that at least one edge in P has load $\Omega(k)$. □

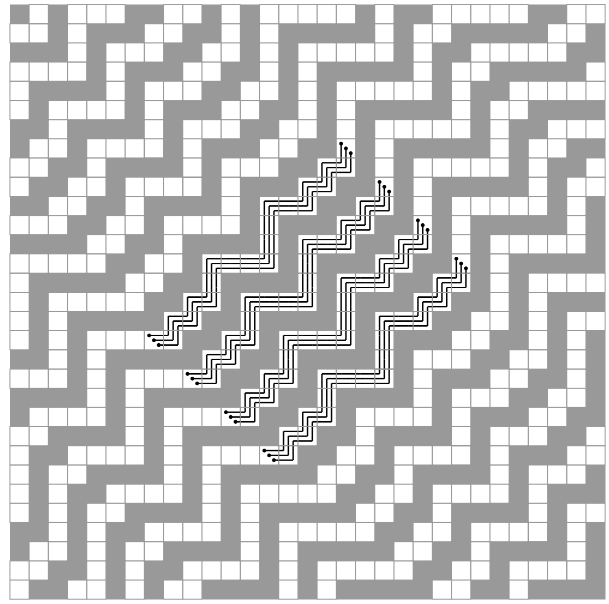

Figure 2: The first level of the grid construction.

5 Semi-oblivious routing in the grid

This section proves that a semi-oblivious routing scheme in the grid which uses a polynomial number of paths must have a competitive ratio bounded below by $\Omega(\log n / \log\log n)$. As was the case with series-parallel graphs, the lower bound comes from constructing a random routing problem in the grid with the same properties (1)-(5) given above in the proof sketch for Section 4. For the grid lower bound, Figure 2 illustrates the random partition of the edge set which is described in step (2), along with the optimal routes for the commodities in class 1. The random stripes running from the lower left to the upper right are parallel copies of a random walk in the grid, each of whose steps is either in the positive x or positive y direction.

A subtlety arises which accounts for the most of the complexity in the proof. In the step labeled as (4) in the proof sketch given above, we must prove that with high probability (over the choice of the random walk path), each path in Π which traverses a constant fraction of the diameter of the grid must have about half of its edges in the black stripes and about half of its edges in the white stripes. Unfortunately, if the stripes were defined using an ordinary (i.e. Markovian) random walk, the walk does not mix rapidly enough to ensure that this happens with sufficiently high probability. Instead we must use a non-Markovian random walk which is designed to have better mixing properties. The non-Markovian walk is defined and analyzed in Section 5.1.

5.1 A non-Markovian random walk on the integers.

Our goal in this section is to define a random walk $\ldots, Z_{-1}, Z_0, Z_1, \ldots$ on the integers which takes steps of bounded size but requires only $O(\ell \operatorname{polylog}(\ell))$ steps to become nearly equidistributed modulo ℓ, for every ℓ. Such a random walk must mix more rapidly than the standard random walk on $\mathbb{Z}$, which requires $\Omega(\ell^2)$ steps to become equidistributed modulo ℓ.

Let $\{X_{i,j} \,:\, i,j \in \mathbb{Z}\}$ be a set of independent random variables, each uniformly distributed on $\{0,1\}$. For an integer $m = 2^{s-1}t$ where s,t are integers and t is odd, let $\eta(m)$ denote the unique element of $\{2^s, 2^s+1, \ldots, 2^{s+1}-1\}$ which is congruent to $\frac{t-1}{2} \pmod{2^s}$. Note that for every integer $k > 1$, the set $\eta^{-1}(k)$ is an arithmetic progression with difference at most k^2. Define a function $\psi : \mathbb{Z} \to \mathbb{Z} \times \mathbb{Z}$ by $\psi(m) = \left(\eta(m), \lfloor m2^{-\eta(m)} \rfloor\right)$. For an integer m let

$$\begin{aligned} Y_m &= X_{\psi(m)} \\ Z_m &= \begin{cases} \sum_{i=1}^{m} Y_i & \text{if } m > 0 \\ 0 & \text{if } m = 0 \\ \sum_{i=1}^{-m} -Y_{-i} & \text{if } m < 0. \end{cases} \end{aligned}$$

Note that $Z_{m+1} - Z_m \in \{0,1\}$ for every integer m.

The following theorem proves the rapid mixing property $(\bmod\ \ell)$ that was asserted earlier. Before stating it, we must define two pieces of notation. First, for any integers $t \leq u$ we let $\mathcal{F}_t$ (resp. $\mathcal{F}_{t,u}$) denote the σ-field generated by the random variables $\{Y_m \,|\, m \leq t\}$ (resp. $\{Y_m \,|\, m \leq t \text{ or } m \geq u\}$). Second, for any positive integer ℓ we define

$$\tau(\ell) = \ell\lceil\log(4\ell^2)\rceil^3.$$

THEOREM 5.1. *Let ℓ be a positive integer and let x be any element of the cyclic group $\mathbb{Z}/(\ell)$. If $u-t > \tau(\ell)$ then*

$$\frac{1}{2\ell} \ \le\ \Pr(Z_u \equiv x \pmod{\ell} \mid \mathcal{F}_t) \ \le\ \frac{3}{2\ell}.$$

Proof. Assume $\ell \geq 2$; otherwise the lemma is trivial. Let h be the function defined on $\mathbb{Z}/(\ell)$ by

$$h(a) = \begin{cases} 1 & \text{if } a \equiv 0 \pmod{\ell} \\ 0 & \text{otherwise.} \end{cases}$$

Let $\omega = e^{2\pi i/\ell}$. Using the identity

$$h(a) = \frac{1}{\ell}\sum_{b=0}^{\ell-1} \omega^{ab},$$

we find that

$$\begin{aligned}\Pr(Z_t \equiv x \pmod{\ell} \mid \mathcal{F}_u) &= \mathbb{E}(h(Z_t - x) \mid \mathcal{F}_u) \\ &= \tfrac{1}{\ell}\textstyle\sum_{b=0}^{\ell-1} \mathbb{E}\left(\omega^{b(Z_t-x)} \mid \mathcal{F}_u\right) \\ &= \tfrac{1}{\ell}\left[1 + \textstyle\sum_{b=1}^{\ell-1} \mathbb{E}\left(\omega^{b(Z_t-x)} \mid \mathcal{F}_u\right)\right] .\end{aligned}$$

To finish, it suffices to prove that $\left|\mathbb{E}\left(\omega^{b(Z_t-x)} \mid \mathcal{F}_u\right)\right| < \frac{1}{2\ell^2}$ for $1 \le b \le \ell - 1$. We may assume without loss of generality that $1 \le b \le \ell/2$, since the numbers $\mathbb{E}\left(\omega^{(\ell-b)(Z_t-x)} \mid \mathcal{F}_u\right)$ and $\mathbb{E}\left(\omega^{b(Z_t-x)} \mid \mathcal{F}_u\right)$ are complex conjugates and hence have the same absolute value.

Let $k \geq 0$ be the least integer such that $2^k b \geq \frac{\ell}{3}$. Observe that $\frac{\ell}{3} \le 2^k b \le \frac{2\ell}{3}$, hence

$$\left|1 + \omega^{2^k b}\right| \le 1.$$

The set $\eta^{-1}(k)$ is an arithmetic progression with difference at most k^2, so the set $J_0 = \eta^{-1}(k) \cap [u+1, t]$ has at least $(t-u)/k^2 > \ell\lceil\log(4\ell^2)\rceil$ elements. Let $J \subseteq J_0$ consist of all elements $m \in J_0$ such that $\lfloor u2^{-k}\rfloor < \lfloor m2^{-k}\rfloor < \lfloor t2^{-k}\rfloor$. The number of elements in the set $J_0 \setminus J$ is at most 2^k, which is less than ℓ, hence $|J| \ge \ell\lceil\log(2\ell^2)\rceil$.

Let $\overline{J} = \{u+1, u+2, \ldots, t\} \setminus J$. Let $\mathcal{F}$ denote the σ-field generated by the random variable $Y = \sum_{j\in\overline{J}} Y_j$ as well as $Z_u, Z_{u-1}, Z_{u-2}, \ldots$. We have

$$Z_t - x = Z_u - x + Y + \sum_{j\in J} Y_j,$$

and the random variables $\{Y_j : j \in J\}$ are independent of $\mathcal{F}$. Moreover, the set $\{Y_j : j \in J\}$ comprises at most 2^k copies of each of the random variables $X(k,i)$, as i runs through the set $I = \{\lfloor j2^{-k}\rfloor \mid j \in J\}$, and these random variables $X(k,i)$ are mutually independent. Hence,

$$\begin{aligned}\mathbb{E}\left(\omega^{b(Z_t-x)} \mid \mathcal{F}\right) &= \omega^{b(Z_u-x+Y)}\mathbb{E}\left(\omega^{\sum_{j\in J} bY_j}\right) \\ &= \omega^{b(Z_u-x+Y)}\mathbb{E}\left(\omega^{\sum_{i\in I} 2^k bX(k,i)}\right) \\ &= \omega^{b(Z_u-x+Y)}\prod_{i\in I}\mathbb{E}\left(\omega^{2^k bX(k,i)}\right) \\ &= \omega^{b(Z_u-x+Y)}\prod_{i\in I}\frac{1}{2}\left(1+\omega^{2^k b}\right) \\ \left|\mathbb{E}\left(\omega^{b(Z_t-x)} \mid \mathcal{F}\right)\right| &= \left|\omega^{b(Z_u-x+Y)}\right|\prod_{i\in I}\frac{1}{2}\left|1+\omega^{2^k b}\right| \\ &\le 2^{-|I|} \le 2^{-|J|/2^k} \le 2^{-|J|/\ell} < \frac{1}{2\ell^2},\end{aligned}$$

as desired. □

COROLLARY 5.1. *Given a sequence $A = \ldots, A_{-1}, A_0, A_1, \ldots$ of residue classes in $\mathbb{Z}/(\ell)$, and a set $R \subseteq \mathbb{Z}/(\ell)$, define $T(A,R)$ to be the random set of integers $T(A,R) = \{s \mid Z_s - A_s \in R\}$. If $u - t = 2m\tau(\ell)$ then*

$$\Pr\left(|T(A,R)\cap[t,u)| < \frac{1}{16}\cdot\frac{|R|}{\ell}\cdot(u-t) \;\middle\|\; \mathcal{F}_t\right) < |R|e^{-m/96}$$

Proof. We will prove the corollary when R is a singleton set $\{r\}$. The full corollary follows using the union bound, together with the observation that $|T(A,R)\cap[t,u)| < \frac{1}{16}\cdot\frac{|R|}{\ell}\cdot(u-t)$ implies that for some $r \in R$, $|T(A,\{r\})| \cap [t,u) < \frac{1}{16\ell}(u-t)$.

For any integer $a \in Z$, let V_a be the random variable which counts the number of s such that $\tau(\ell) \le s - a < 2\tau(\ell)$ and $s \in T(A,\{r\})$. Let $W_a = \min\{V_a, 6\tau(\ell)/\ell\}$. From Theorem 5.1 we have

$$\mathbb{E}(W_a \mid \mathcal{F}_a) \ge \frac{1}{4}\tau(\ell)/\ell. \tag{5.1}$$

Let $s = 2\tau(\ell)$. Equation (5.1) establishes that the sequence of random variables $\widetilde{W}_i = V_{u+is} - \frac{1}{4}\tau(\ell)/\ell$ are a martingale difference sequence. The differences are bounded by $6\tau(\ell)/\ell$. Applying Azuma's inequality for submartingales, we find that

$$\begin{aligned}&\Pr\left(\sum_{i=0}^{m-1} W_{u+is} < \frac{m}{8}\tau(\ell)/\ell \;\middle\|\; \mathcal{F}_t\right) \\ <\ &\exp\left(-\left(\frac{m\tau(\ell)}{8\ell}\right)^2\left(\frac{1}{12m(\tau(\ell)/\ell)^2}\right)\right) = e^{-m/96}.\end{aligned}$$

To conclude the proof, we observe that $\sum_{i=0}^{m-1} W_{u+is} \le \sum_{i=0}^{m-1} V_{u+is} \le |T(A,\{r\})\cap[t,u)|$, and that $\frac{m}{8}\tau(\ell)/\ell = \frac{1}{16\ell}(u-t)$. □

5.2 Random height functions and routing problems in the infinite grid. We may define a random function $H : \mathbb{Z} \times \mathbb{Z} \to \mathbb{Z}$ using the random walk $(Z_i)_{i \in \mathbb{Z}}$ defined in Section 5.1, as follows:

$$H(x, y) = Z_{x+y} - x.$$

This function has the following two properties:

1. For any integers a, x, y, $H(x-a, y+a) = H(x, y)+a$.
2. For any integers x, y, exactly one of the relations $H(x+1, y) = H(x, y)$, $H(x, y+1) = H(x, y)$ holds.

Property (1) implies that the function $(x, y) \mapsto (H(x, y), x+y)$ is invertible. In fact, its inverse is the function

$$\Upsilon(z, w) = (H(0, w) - z, w + z - H(0, w)).$$

Property (2) implies that the sequence $P_z = (\ldots, \Upsilon(z, -1), \Upsilon(z, 0), \Upsilon(z, 1), \ldots)$ is an infinite path in $\mathbb{Z} \times \mathbb{Z}$ such that each point is joined to the next by the vector $(1, 0)$ or $(0, 1)$. (These infinite paths are the random walk paths represented by the black and white stripes in Figure 2.)

For an integer $k \geq 2$, let $\nu_k(x)$ denote the unique integer in the interval $\left(-\frac{k^5}{2}, \frac{k^5}{2}\right]$ which is congruent to $\lfloor \frac{1}{2} + \frac{x}{((k-1)!)^5} \rfloor$ modulo k^5. For an integer x let $\lambda(x)$ denote the largest k such that $\nu_k(x)$ is odd, or $\lambda(x) = 1$ if no such k exists. (There are always only finitely many k such that $\nu_k(x)$ is odd, because $\nu_k(x) = 0$ whenever $(k-1)! > |2x|^{1/5}$.) For a point $s = (x, y) \in \mathbb{Z}^2$, let $L(s) = (\lambda(H(s))!)^5$. Define an infinite demand matrix $D = (D_{st})_{s,t \in \mathbb{Z}}$ as follows. For a point $s = (x, y)$, put $t(s) = \Upsilon(H(s), x + y + L(s))$ and let

$$D_{st} = \begin{cases} 1/L(s) & \text{if } t = t(s) \\ 0 & \text{otherwise.} \end{cases}$$

LEMMA 5.1. *There exists a flow* $\mathbf{f}$ *in the infinite grid which satisfies* D *and whose edge congestion is equal to* 1.

Proof. For a vertex s, let $z = H(s)$. Recall the path $P_z = (\ldots, \Upsilon(z, -1), \Upsilon(z, 0), \Upsilon(z, 1), \ldots)$ defined earlier. Define $P(s)$ to be the subpath of P_z joining s to $t(s)$. The flow f sends $1/L(s)$ units of flow on path $P(s)$, for every $s \in \mathbb{Z}^2$. For an edge e of the grid, either e does not belong to any of the paths P_z $(z \in \mathbb{Z})$ and f sends no flow on e, or e belongs to exactly one path P_z and there are $(\lambda(z)!)^5$ paths $P(s) \subset P_z$ containing e, each of which carries $(\lambda(z)!)^{-5}$ units of flow. In the former case, the congestion of e is 0; in the latter case the congestion of e is 1. □

The next lemma concerns the function

$$\lambda_k(x) = \begin{cases} \lambda(x) & \text{if } 2 \leq \lambda(x) \leq k \\ 1 & \text{otherwise.} \end{cases}$$

LEMMA 5.2. *Fix some* $k \geq 2$ *and let* $m = (k!)^5$. *The values of* $\lambda_k(x)$ *and of* $\nu_i(x)$ $(2 \leq i \leq k)$ *depend only on the residue class of* x *modulo* m. *If* x *is a uniformly random element of* $\mathbb{Z}/(m)$ *then the random variables* $\nu_i(x)$ *are mutually independent, and each is uniformly distributed in its range. For* $1 \leq j \leq k$, *the number of elements* $x \in \mathbb{Z}/(m)$ *which satisfy* $\lambda_k(x) = j$ *is at most* $4 \cdot 2^{j-k} m$. *If* $i \geq j$ *and* I *is a subinterval of* $\left(-\frac{i^5}{2}, \frac{i^5}{2}\right]$ *whose length is denoted by* $|I|$, *then the number of* $x \in \mathbb{Z}/(m)$ *which satisfy* $\lambda_k(x) = j$ *and* $\nu_i(x) \in I$ *is at least* $\left(\frac{|I|-2}{4i^5}\right) 2^{j-k} m$.

Proof. If $x \equiv y \pmod{m}$ then the numbers $\frac{1}{2} + \frac{x}{((i-1)!)^5}$ and $\frac{1}{2} + \frac{y}{((i-1)!)^5}$ differ by a multiple of i^5 for every $i \leq k$. Consequently, $\nu_i(x) = \nu_i(y)$ for every $i \leq k$. The value of $\lambda_k(x)$ is completely determined by the values of $\nu_2(x), \nu_3(x), \ldots, \nu_k(x)$, hence the residue class of x modulo m determines $\lambda_k(x)$ as well as $\nu_2(x), \ldots, \nu_k(x)$.

Let $m' = ((k-1)!)^5$, and let h be the homomorphism $\mathbb{Z}/(m) \to \mathbb{Z}/(m')$ induced by the canonical homomorphism $\mathbb{Z} \to \mathbb{Z}/(m')$. For a given element $y \in \mathbb{Z}/(m')$, the set $h^{-1}(y)$ has exactly k^5 elements, and as x runs through the elements of $h^{-1}(y)$, the function $\nu_k(x)$ takes each integer value in the interval $\left(-\frac{k^5}{2}, \frac{k^5}{2}\right]$ exactly only. It follows that if x is a uniformly random residue class in $\mathbb{Z}/(m)$, the random variable $\nu_k(x)$ is uniformly distributed in $\mathbb{Z} \cap \left(-\frac{k^5}{2}, \frac{k^5}{2}\right]$, and is independent of $h(x)$. From the first paragraph of this proof, we know that $h(x)$ determines the values of $\nu_i(x)$ for $2 \leq i < k$. By induction on k, we may conclude that the random variables $\nu_i(x)$ $(2 \leq i \leq k)$ are mutually independent, and each is uniformly distributed in its range.

In particular, let $E_i, F_i, G_i(I)$ respectively denote the events that $\nu_i(x)$ is even, that $\nu_i(x)$ is odd, and that $\nu_i(x)$ is an element of the interval I which is odd if $i > 1$, even if $i = 1$. We have

$$\begin{array}{ccc} \frac{1}{2} - \frac{1}{i^5} & \leq \Pr(E_i) \leq & \frac{1}{2} + \frac{1}{i^5} \\ \frac{1}{2} - \frac{1}{i^5} & \leq \Pr(F_i) \leq & \frac{1}{2} + \frac{1}{i^5} \\ \frac{|I|-2}{2i^5} & \leq \Pr(G_i(I)) \leq & \frac{|I|+2}{2i^5}. \end{array}$$

All the remaining claims in the statement of the lemma follow from these probability estimates, along with the fact that these events are mutually independent for different values of i. □

5.3 Proof of the lower bound. Let the random function H and the random demand matrix $(D_{st})_{(s,t) \in \mathbb{Z}^2}$ be defined as in the preceding section. Let $\Lambda \subset \mathbb{Z}^2$ denote the set of all (x, y) such that $\lambda(H(x, y)) = 1$. Similarly, let $\Lambda_k \subset \mathbb{Z}^2$ denote the set of all (x, y) such that $\lambda_k(H(x, y)) = 1$.

LEMMA 5.3. *Suppose $\lambda(H(s)) = k$ and $\nu_k(H(s)) \in \left(-\frac{k^5}{6}, \frac{k^5}{6}\right]$. Suppose moreover that k is even. For any path P from s to $t(s)$, the probability that $|P \cap \Lambda| < 2^{-k-4}L(s)$ is less than $e^{-k^{1.5}}$ provided that k is sufficiently large.*

Proof. We focus first on the probability $\Pr(|P \cap \Lambda_k| < 2^{-k}k^{-4}L(s))$. Write $s = (x_0, y_0)$ and let $w = x_0 + y_0$. For $j = 1, \ldots, L(s)$ let $v_j = (x_j, y_j)$ denote the first vertex in P to satisfy $x_j + y_j = w + j$. (Such a vertex exists, for each i, because as one traverses P the parameter $x + y$ starts at w and finishes at $w + L(s)$, and it increases by at most 1 on each step.) We have $H(v_j) = Z_{w+j} - x_j$. We shall apply Corollary 5.1 with

$$A_i = \begin{cases} x_{i-w} & \text{if } w \le i \le w + L(s) \\ 0 & \text{otherwise.} \end{cases}$$

Let $\ell = 2((k-1)!)^5$. Since k is even, for any integer z the values of $\nu_2(z), \ldots, \nu_k(z)$ — and therefore also of $\lambda_k(z)$ — are determined by the residue class of z in $\mathbb{Z}/(\ell)$. Let R denote the set of residue classes $z \in \mathbb{Z}/(\ell)$ such that $\lambda_k(z) = 1$. Note that $|R| = 2^{-k}\ell$, and that $(x, y) \in \Lambda_k$ if and only if there exists some $r \in R$ such that $H(x, y) \equiv r \pmod{\ell}$. Defining the sequence A as above, we see that

$$|P \cap \Lambda_k| \ge |T(R, A) \cap [w, w + L(s)]|.$$

Now, recall that $\tau(\ell) = \ell\lceil\log(4\ell^2)\rceil^3$ which is bounded above by $16((k-1)!)^5(k \log k)^3$. So $L(s) > \frac{1}{16}\tau(\ell)k^2(\log k)^{-3}$, which is greater than $192k^{1.5}\tau(\ell)$ for sufficiently large k. Applying Corollary 5.1 with $m = 96k^{1.5}$ we conclude that

$$\Pr(|P \cap \Lambda_k| < 2^{-k-4}L(s)) < e^{-k^{1.5}}.$$

To finish, we will prove that if $|P \cap \Lambda_k| > |P \cap \Lambda|$ then $|P \cap \Lambda| \ge 2^{-k}k^{-4}L(s)$. Indeed, if P intersects $\Lambda_k \setminus \Lambda$ it means that $\lambda(H(v)) = j > k$ for some $v \in P$. In turn, this implies that $\nu_j(H(v))$ is odd. But by our hypothesis on s, $\nu_j(x) = 0$ for every x such that $|x - s| < (k!)^5/6$, so we may conclude that $|H(v) - H(s)| \ge (k!)^5/6$. As we move along the portion of P which joins s to v, the function H assumes every value between $H(s)$ and $H(v)$. By Lemma 5.2, at least $(1 - 12/k^5)2^{-k}m/12$ of these values belong to $\lambda^{-1}(1)$. Hence $|P \cap \Lambda| \ge (1 - 12/k^5)2^{-k}m/12$, which is greater than $2^{-k-4}m$ for sufficiently large k. □

THEOREM 5.2. *For any positive integer d, if $n > ((2d)!)^5$ and Π is a set of fewer than n^d paths in the $n \times n$ grid $\mathbb{G}_n$, then there is a demand matrix D for $\mathbb{G}_n$ which can be satisfied by a multicommodity flow of edge congestion 1, but any Π-valued flow satisfying D has edge congestion $\Omega(\log n / \log\log n)$.*

Proof. Sample a random demand matrix D from the distribution defined in the preceding section, and keep only those commodities whose source and sink both belong to the set $\{1, 2, \ldots, n\}^2$. In this way we obtain a random demand matrix for $G = \mathbb{G}_n$. The fact that D can be satisfied by a flow of edge congestion 1 now follows directly from Lemma 5.1. We will prove that with high probability, every Π-valued flow satisfying D has edge congestion $\Omega(\log n / \log\log n)$. For a pair of vertices s, t, let $d_\Pi(s, t)$ denote the minimum of $|P \cap \Lambda|$ over all paths $P \in \Pi$ joining s to t. Let k be the largest integer such that $(k!)^5 < n$. We will show that

1. $|V(G) \cap \Lambda| = O\left(2^{-k}n^2\right)$.
2. $\sum_{s,t} D_{st}d_\Pi(s, t) = \Omega\left(2^{-k}n^2 \log n / \log\log n\right)$.

From these two facts, it follows immediately that for any Π-valued flow satisfying D, there is a vertex $v \in \Lambda$ whose congestion is $\Omega(\log n / \log\log n)$. Since the edge congestion of a flow in G is at least one-fourth of the vertex congestion, this proves the theorem.

The first fact follows easily from Lemma 5.2. To prove the second fact, we observe that the height function H maps $V(G)$ to a set of n consecutive integers which includes 0. It is easy to prove that the set of all h such that $|H^{-1}(h) \cap V(G)| > n/2$ contains a subset X consisting of at least $n/4$ consecutive integers. Let X_j be the subset consisting of all $x \in X$ such that $\lambda(x) = j$ and $\nu_j(x) \in \left(-\frac{j^5}{6}, \frac{j^5}{6}\right]$. By Lemma 5.2 we obtain the lower bound $|X_j| = \Omega(2^{j-k}n)$. Let $S_j = H^{-1}(X_j)$; note that by construction, $|S_j| = \Omega(2^{j-k}n^2)$. If $((d+4)\log n)^{2/3} < j \le k$ and j is even, we may apply Lemma 5.3 to say, for every path $P \in \Pi$ which joins s to $t(s)$ for some $s \in S_j$, that with probability at least $1 - 1/n^{d+3}$ the set $P \cap \Lambda$ has $\Omega(2^{-j}(j!)^5)$ vertices. Taking a union bound over all such paths P, we may say that with probability $1 - 1/n$ we have $d_\Pi(s, t(s)) = \Omega(2^{-j}(j!)^5)$ for all even j between $j_0 = ((d+4)\log n)^{2/3}$ and k and for all $s \in S_j$. Recalling that $D_{s,t(s)} = (j!)^{-5}$, we obtain

$$\begin{aligned}
&\sum_{s,t} D_{st}d_\Pi(s, t(s)) \\
&\quad \ge \sum_{j_0 < j \le k, j \text{ even}} \ \sum_{s \in S_j} D_{s,t(s)}d_\Pi(s, t(s)) \\
&\quad \ge \sum_{j_0 < j \le k, j \text{ even}} \Omega(2^{j-k}n^2)(j!)^{-5}\Omega(2^{-j}(j!)^5) \\
&\quad = \Omega(2^{-k}n^2k) = \Omega(2^{-k}n^2 \log n / \log\log n).
\end{aligned}$$

□

6 Lower bounds for the number of bends

In this section, we prove Theorem 3.4 by showing that in a grid it is impossible to route multicommodity flows with constant (or even polylogarithmic) competitive ratio for

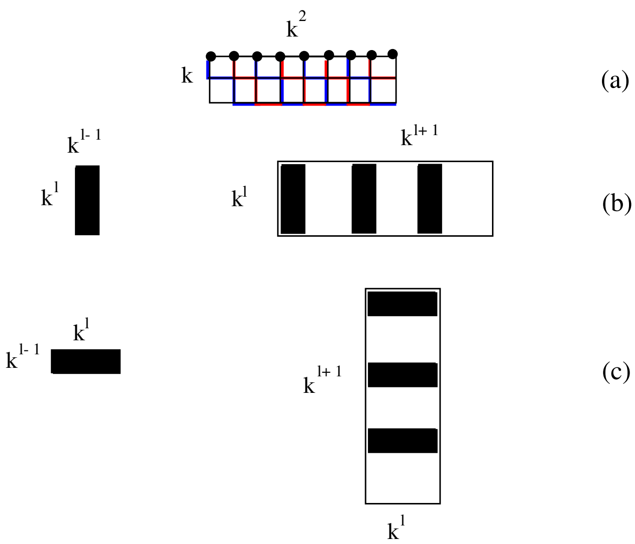

Figure 3: Construction illustration for Theorem 3.4 (a) building block (level one) (b) construction for odd level numbers (c) construction for even level numbers.

congestion while having only a constant number of bends per each flow path.

We start with a level one ($\ell = 1$) building block as follows. We consider a k-by-k^2 grid, in which there is a source in each node at the first row (thus there are k^2 sources in total). We will determine k in terms of n later in the proof. Now we extend this construction for levels $l > 1$ as follows (See Figure 3.) For level $\ell = 2\ell'$ ($\ell = 2\ell' + 1$), we construct building blocks of size $k^{\ell+1}$-by-k^{ℓ} (k^{ℓ}-by-$k^{\ell+1}$) by taking k blocks of level $\ell - 1$ and place them on the same set of columns (rows) with $k^{\ell} - k^{\ell-1}$ row (column) spacing between the consecutive blocks and after the last block. We define our graph to be a block of level $2b$ with $k = n^{\frac{1}{2b+1}}$ (add some dummy columns before the first column to make the graph n-by-n grid). We will put the sources on some vertices of the last column and the last row as follows. Each source s has a unit demand for the unique sink t on the last column or the last row which lies on the same diagonal that source s lies.

First we observe that there is one source per diagonal because of spacing that we used to create the blocks at each stage. Thus we can route the flow from each source s with edge-congestion one by taking repetitively unit vertical and unit horizontal edges along the 45 degree diagonal (see Figure 3.a) until we reach to the unique node t at the last column or the last row of the block of level i, for $1 \leq i \leq 2b$. Thus we can satisfy all the unit demands betweens sources and their corresponding sinks with edge-congestion one.

In a block of level $\ell = 2\ell' - 1$, there are $k^{\ell+1}$ sources and at least $2/3$ of them have to exit horizontally (otherwise, we have congestion more than $k/6$ in the columns). In a $\ell = 2\ell'$, there are $k^{\ell+1}$ sources and $2/3$ of them have to exit vertically. Thus $2/3$ of the flows have to run horizontally at the odd stages and $2/3$ have to run vertically at the even stages. This means that $1/3$ of the flows have to bend in each pair of levels. Hence, the average number of bends over all levels and flow paths is at least b. Thus there is a flow which bends at least b time which is a contradiction to our assumption.

7 Node-capacitated and directed graphs

In this section, we prove Theorem 3.5 by constructing a node-capacitated undirected graph as a counterexample (we can make this an edge-capacitated directed graph by standard reductions). The counterexample is a graph with a single sink node w attached to the right side of a bipartite graph G. The graph G is constructed in [3]. We repeat the definition of G here, and we recall its basic properties. Let d be a positive integer, and let X be the ring $(\mathbb{F}_2)^d$ (i.e. the cartesian product of d copies of the fields $\mathbb{F}_2 = \{0, 1\}$, with addition and multiplication defined componentwise). Let $b = 1 - \frac{1}{2}\log_2(3)$. Considering X as a vector space over $\mathbb{F}_2$, let Y be a linear subspace of dimension $\lfloor bd \rfloor$ such that the cardinality of the set $Z = \{z \in X \mid zy \neq 0 \text{ for all nonzero } y \in Y\}$ is as large as possible. G is a bipartite graph $(V_L \cup V_R, E)$ where $V_L = X \times Z$, $V_R = X$, and $E = \{((x_L, z_L), x_R) \mid x_R - x_L = yz_L \text{ for some } y \in Y\}$.

Letting $n = |V(G)|$, the following properties of G are proven in [3]:

1. $n = 4^{d-o(d)}$.
2. Every vertex in V_L has degree $2^{bd} = n^{0.2075\ldots}$.
3. The edge set of G is partitioned into matchings $\{M_i\}_{i \in I}$, each of size $2^{d-o(d)} = n^{0.4999\ldots}$.
4. Each M_i is an induced matching, i.e. if M_L, M_R denote the sets of left and right endpoints of edges of M_i, respectively, then there is no edge from M_L to M_R except for the edges of M_i.
5. For each such matching M_i, if $\Gamma(M_L)$ denotes the set of all vertices adjacent to M_L, then $|\Gamma(M_L) \setminus M_R| < 3^{d/2} = n^{0.39624\ldots}$.

Let $\widetilde{G}$ be the graph obtained from G by adjoining a single sink vertex w with edges (v, w) for every $w \in V_R$. Suppose Π is a semioblivious routing scheme for $\widetilde{G}$ such that $|\Pi(s, t)| = O(n^{1/5})$ for every pair s, t. We define a demand matrix $D(i)$ for each matching $M = M_i$ ($i \in I$) by specifying that $D(i)_{s,t} = 1$ if $s \in M_L, t = w$, otherwise $D(i)_{s,t} = 0$. Choose the matching M uniformly at random. For every $s \in V_L$, the neighbor of s in M conditional on the event $s \in M_L$ is uniformly distributed among the $n^{0.2075\cdots}$ neighbors of s. (This follows from the fact that the matchings M_i are a partition of the edge set E.) Since there are only $n^{0.2}$ paths in $\Pi(s,w)$, the probability that $\Pi(s,w)$ contains a path starting with an edge of M is only $n^{-0.0075}$, so the expected number of vertices s that have such a path is $|M|^{-0.0075}$. Let us now fix a matching $M = M_i$ $(i \in I)$ such that fewer than $|M|/2$ vertices $s \in V_L$ have a path in $\Pi(s,w)$ starting with an edge of M. Recall that M is an *induced* matching, so if $s \in M_L$ and P is a path from s to w which does not begin with an edge of M, then P passes through a vertex in $\Gamma(M_L) \setminus M_R$. If we satisfy the demand in $D(i)$ using a Π-valued flow, then at least $|M|/2 = n^{0.4999\cdots}$ units of flow must therefore go through the vertex set $\Gamma(M_L) \setminus M_R$. This vertex set has cardinality $n^{0.39624\cdots}$ so at least one of its vertices has congestion greater than $n^{1/10}$.

Acknowledgments

We thank Harald Räcke for helpful discussions.

References

[1] M. ANDREWS, D. BHATIA, T. LEIGHTON, F. MAKEDON, C. H. NORTON, AND Y. L. ZHANG, *Improved algorithms for routing on two-dimensional grids*. Unpublished manuscript available at *citeseer.ist.psu.edu/249177.html*.

[2] D. O. AWDUCHE, *MPLS and traffi c engineering in ip networks*, IEEE Communications Magazine, (1999), pp. 42–47.

[3] B. AWERBUCH, M. HAJIAGHAYI, R. KLEINBERG, AND T. LEIGHTON, *Online client-server load balancing without global information*, in Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms, ACM Press, 2005, pp. 197–206.

[4] Y. AZAR, E. COHEN, A. FIAT, H. KAPLAN, AND H. RÄCKE, *Optimal oblivious routing in polynomial time*, in Proceedings of the 35th ACM symposium on Theory of Computing, ACM Press, 2003, pp. 383–388.

[5] M. BIENKOWSKI, M. KORZENIOWSKI, AND H. RÄCKE, *A practical algorithm for constructing oblivious routing schemes*, in Proceedings of the 15th ACM Symposium on Parallel Algorithms and Architectures, ACM Press, 2003, pp. 24–33.

[6] C. HAIBT-NORTON, *Promlems in Discrete Optimization*, PhD thesis, Massachusetts Institute of Technology, Cambridge, MA 02139, USA, 1993.

[7] M. HAJIAGHAYI, R. D. KLEINBERG, T. LEIGHTON, AND H. RÄCKE, *Oblivious routing on node-capacitated and directed graphs*, in Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms, ACM Press, 2005, pp. 782–790.

[8] C. HARRELSON, K. HILDRUM, AND S. RAO, *A polynomial-time tree decomposition to minimize congestion*, in Proceedings of the 15th Annual ACM Symposium on Parallel Algorithms and Architectures, 2003, pp. 34–43.

[9] R. KARP, T. LEIGHTON, R. RIVEST, C. THOMPSON, U. VAZIRANI, AND V. VAZIRANI, *Global wire routing in two-dimensional arrays*, Algorithmica, 2 (1987), pp. 113–129.

[10] T. LEIGHTON, *Research Topics in Theory of Parallel Computation and VLSI*. Course Notes from Fall 1985.

[11] H. RÄCKE, *Minimizing congestion in general networks*, in Proceedings of the 43rd Symposium on Foundations of Computer Science, IEEE Computer Society, 2002, pp. 43–52.