

Balancing Traffic Load Using One-Turn Rectilinear Routing

Stephane Durocher¹, Evangelos Kranakis², Danny Krizanc³, and Lata Narayanan⁴

¹ School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada
`sdurocher@cs.uwaterloo.ca`

² School of Computer Science, Carleton University, Ottawa, Ontario, Canada
`kranakis@scs.carleton.ca`

³ Department of Mathematics and Computer Science, Wesleyan University, Middletown, Connecticut, USA `dkrizanc@wesleyan.edu`

⁴ Department of Computer Science, Concordia University, Montréal, Québec, Canada `lata@cse.concordia.ca`

Abstract. We consider the problem of load-balanced routing, where a dense network is modelled by a continuous square region and origin and destination nodes correspond to pairs of points in that region. The objective is to define a routing policy that assigns a continuous path to each origin-destination pair while minimizing the traffic, or *load*, passing through any single point. While the average load is minimized by straight-line routing, such a routing policy distributes the load non-uniformly, resulting in higher load near the center of the region. We consider one-turn rectilinear routing policies that divert traffic away from regions of heavier load, resulting in up to a 33% reduction in the maximum load while simultaneously increasing the path lengths by an average of less than 28%. Our policies are simple to implement, being both local and oblivious. We provide a lower bound that shows that no one-turn rectilinear routing policy can reduce the maximum load by more than 39% and we give a polynomial-time procedure for approximating the optimal randomized policy.

1 Introduction

The problem of routing in multi-hop wireless networks has received extensive attention in the last decade [1, 2, 12, 14, 18]. Many of the proposed routing protocols attempt to find shortest paths between pairs of nodes, or try to bound the *stretch factor* of the paths, while trying to ensure that the paths are loop-free. This approach takes into account a single packet traversing the network and tries to optimize performance for this packet. A more global and realistic view would consider the performance of the protocol under the assumption of many traffic flows in the network. In this situation, there can often be *congestion* created by several packets that need to be forwarded by the same intermediate nodes at the same time. This congestion is very likely to influence the latency experienced by

a packet. A routing protocol should therefore attempt to avoid creating highly congested nodes. Not only does this improve packet latency, it would also improve the lifetime of a wireless network, where heavily loaded nodes may run out of battery power and disconnect the network.

In this paper, we investigate routing protocols for wireless networks with the aim of minimizing the congestion experienced at nodes. We consider a multi-hop ad hoc network consisting of identical location-aware nodes, uniformly and densely deployed within a given planar region. Furthermore, we assume that the traffic pattern is uniform point-to-point communication, i.e., each node has the same number of packets to send to every other node in the network. This is sometimes called the *all-to-all* communication pattern. A routing policy must define, for every ordered pair of nodes (u, v) , a path in the network to get from u to v . The *load* at a given node v is the number of paths that pass through v . The average (maximum) load for a network with respect to a particular routing policy is the average (respectively maximum) load over all nodes in the network. The fundamental question we wish to answer is: *what routing policy minimizes the maximum load in the network?*

It seems intuitively evident that for nodes within a convex planar region, shortest path routing should cause maximum load near the *geometric center*. Indeed, this has been proved analytically for disks (see for example [16]) and squares and rectangles (see Section 3). This suggests that if load balancing is a fundamental concern then a good routing policy should redirect some of the traffic away from the geometric center and other areas of high load. However, load balancing cannot be the *only* concern: taking unnecessarily long paths just to bypass the center can drastically increase the stretch factor and the average load of nodes in the network, and can therefore be very inefficient in terms of energy consumption. Furthermore, it is critical that the forwarding strategy required to implement the routing policy be simple and have low memory requirements. Ideally, the routing policy should be *oblivious* (the route between u and v depends only on the identities or locations of u and v) and the forwarding strategy should be *local* (the forwarding node can make its decision based only on itself and its neighbors, and the packet header contains only the address of the destination).

In the setting of nodes uniformly distributed in a given planar convex region, very little research has been done on finding a simple routing policy that achieves both a reasonable stretch factor and a minimum value of maximum congestion. In [6], an algorithm achieving a good tradeoff between stretch factor and load balance is shown for the special case when all nodes are located in a narrow strip of width at most 0.86 times the transmission radius. The analysis is not specific to the all-to-all communication pattern. Popa *et al.* [16] address the all-to-all routing problem for the case when the region containing the nodes is a unit disk. They establish quantitatively the *crowded center* effect for shortest-path routing as a nearly-quadratic function that peaks at the center of the disk and present a theoretical approach that is guaranteed to find paths minimizing the maximum load. They also give a practical solution (curveball routing) whose performance

compares favorably to the optimum. No theoretical bounds are given on the stretch factor of the routes for either strategy.

In this paper, we investigate the problem of load-balanced routing when the nodes are uniformly and densely packed in a square or rectangular region. As in [16], our approach is to look at the unit square (and the $k \times 1$ rectangle) as a *continuous space* rather than formed by discrete nodes. This makes it possible to analyze the average and maximum load induced by a routing policy, without regard to the topology of the actual network. At the same time, the results should predict the behavior of a network with very densely and uniformly deployed nodes. Shortest-path routing corresponds to *straight-line routing* in this setting. We derive the average and maximum load for straight-line routing in a unit square and confirm the crowded-center effect for squares and rectangles. In keeping with the goal of minimizing congestion while ensuring a reasonable stretch factor, we investigate the class of *rectilinear* routing policies that assign to each origin-destination pair of nodes one of the two possible rectilinear paths containing only one turn. It is not difficult to show that all such one-turn rectilinear strategies have a maximum stretch factor of $\sqrt{2}$. Furthermore, they are simple and realistic in the ad hoc network setting; the routing policy is oblivious and the forwarding algorithm is local. We propose and analyze several simple rectilinear strategies, the best of which reduces the maximum load by about 33% compared to the straight-line policy. We also characterize the optimal randomized rectilinear policy as the solution to an optimization problem and provide an efficient procedure for approximating it.

1.1 Overview of Results

Our main contributions are summarized below:

- We derive an exact expression for the load induced by a straight-line routing policy at an arbitrary point in the unit square. We show that the average and maximum load for the straight-line routing policy are 0.5214 and 1.1478 respectively.
- We show that the average load for *every* one-turn rectilinear routing policy is $2/3$. The maximum and average stretch factor for such policies are shown to be $\sqrt{2}$ and 1.2737 respectively.
- We propose several one-turn rectilinear routing policies and derive their maximum load. The best of these, called the *diagonal rectilinear policy*, achieves a maximum load of 0.7771, which represents a 33% improvement over straight-line routing.
- We prove a lower bound of 0.7076 on the the maximum load for *any* one-turn rectilinear policy.
- We characterize the optimal randomized rectilinear policy as the solution to an optimization problem and provide an efficient procedure for approximating it. Numerical results suggest that the maximum load for the best possible rectilinear policy is close to 0.74.

Detailed proofs for some results are omitted due to space restrictions.

1.2 Related work

In this section, we briefly describe other efforts to address the congestion problem. Several studies confirm the crowded center effect for shortest path routing [8, 11, 15, 16]. In [15, 16], the load at the center of a circular area is derived analytically, by modelling the area as a continuous region, as in this paper, rather than as formed by discrete nodes. The node distribution resulting from a random waypoint mobility model in an arbitrary convex domain is analyzed in [10]; this is related to the load probability density for straight-line routing.

The tradeoffs between congestion and stretch factor in wireless networks has been studied in [13] and [7]. For instance, for growth-bounded wireless networks, Gao and Zhang [7] show routing algorithms that simultaneously achieve a stretch factor of c and a load balancing ratio of $O((n/c)^{1-1/k})$ where k is the growth rate. (The load balancing ratio is defined to be the ratio between the maximum load on any node induced by the algorithm versus that created by the optimal algorithm.) They also derive an algorithm for unit disk graphs with bounded density and show that if the density is constant, shortest path routing has a load balancing ratio of $\Theta(\sqrt{n})$. The communication patterns considered are arbitrary, the lower bound does not derive from the all-to-all communication pattern, and the routing algorithms are not oblivious.

The all-to-all communication pattern has been studied extensively in the context of interconnection networks, and particularly in WDM optical networks. In this context, [4] defined the forwarding index of a communication network with respect to a specific routing algorithm to be the maximum number of paths going through any vertex in the graph. The forwarding index of the network itself is the minimum over all possible routing algorithms for the network. This notion was extended to the maximum load on an edge [9], which is more appropriate to wired networks. However, for wireless networks, the node forwarding index captures the load on a wireless node better. While the node forwarding index for specific networks, including the ring and torus networks has been derived exactly [4], it has not been studied for two-dimensional grid networks, which would perhaps be a good approximation for the dense wireless networks of interest to us. Our results in Section 6 provide an approximation for the forwarding index in grid graphs for the class of one-turn rectilinear routing schemes.

There does not appear to be much work on routing with a view to reducing the congestion for the all-to-all communication pattern in specific planar regions, the model of interest in this paper. As stated earlier, [6] looks at nodes contained in a narrow strip and [16] addresses the problem for the unit disk. Busch et al. [3] analyze routing on embedded graphs via a random intermediate point located near the perpendicular bisector of the origin and destination; we consider the generalization of this strategy to convex regions in Section 4.4. Popa *et al.* [16] give expressions for the maximum and average load induced by straight-line routing in unit disks, and propose a practical algorithm called curveball routing whose performance is close to the optimum for disks. They also provide experimental results on greedy routing versus curveball routing in square- and

rectangular-shaped areas, and show that curveball routing achieves a reduction in load in such areas, but they do not provide any theoretical results.

2 Definitions

2.1 Routing Policies and Traffic Load

Given a convex region $A \subseteq \mathbb{R}^2$, a *routing policy* P assigns a route to every origin-destination pair $(u, v) \in A^2$, where the *route* from u to v , denoted $\text{route}_P(u, v)$, is a plane curve segment contained in A , whose endpoints are u and v . For a given routing policy P on a region A , the traffic load at a point p is proportional to the number of routes that pass through p . Formally,

Definition 1. *Given a routing policy P on a region A , the load at point p is*

$$\lambda_P(p) = \iint_A f_P(p, u, v) du dv, \quad \text{where } f_P(p, u, v) = \begin{cases} 1 & \text{if } p \in \text{route}_P(u, v), \\ 0 & \text{otherwise.} \end{cases}$$

The *average load* of routing policy P on region A is given by

$$\lambda_{\text{avg}}(P) = \frac{1}{\text{Area}(A)} \int_A \lambda_P(p) dp, \quad (1)$$

where $\text{Area}(A) = \int_A dp$ denotes the area of region A . The average length of a route determined by policy P between two points in A is given by

$$\text{length}_{\text{avg}}(P) = \frac{1}{\text{Area}(A)^2} \iint_A \text{length}(\text{route}_P(p, q)) dq dp. \quad (2)$$

Since $\text{length}(\text{route}_P(u, v)) = \int_A f_P(p, u, v) dp$, Proposition 1 follows from (1) and (2):

Proposition 1. *Given routing policy P on a region A ,*

$$\lambda_{\text{avg}}(P) = \text{Area}(A) \cdot \text{length}_{\text{avg}}(P). \quad (3)$$

In addition to average load, a routing policy P on a region A is also characterized by its *maximum load*, given by

$$\lambda_{\text{max}}(P) = \max_{p \in A} \lambda_P(p). \quad (4)$$

2.2 Straight-Line Routing Policy

The *straight-line routing policy*, denoted S , assigns to every pair (u, v) the route consisting of the line segment between u and v . In straight-line routing,

$$\text{length}(\text{route}_S(p, q)) = \|p - q\| = \sqrt{(p_x - q_x)^2 + (p_y - q_y)^2}. \quad (5)$$

Since the line segment from u to v is the shortest route from u to v , it follows that straight-line routing minimizes (2). Consequently, for any convex region A and any routing policy $P \neq S$,

$$\lambda_{\text{avg}}(S) \leq \lambda_{\text{avg}}(P). \quad (6)$$

The *average stretch factor* and *maximum stretch factor* of routing policy P on region A are respectively given by

$$\text{str}_{\text{avg}}(P) = \frac{1}{\text{Area}(A)^2} \iint_A \frac{\text{length}(\text{route}_P(p, q))}{\text{length}(\text{route}_S(p, q))} dq dp \quad (7)$$

$$\text{str}_{\text{max}}(P) = \max_{\{p, q\} \subseteq A} \frac{\text{length}(\text{route}_P(p, q))}{\text{length}(\text{route}_S(p, q))}. \quad (8)$$

2.3 One-Turn Rectilinear Routing Policies

Recent related work on this problem has considered the case when region A is a disk [16]. In this paper, we consider the case when region A is bounded by a square or a rectangle. As we show in Section 3, the load in straight-line routing on a square or a rectangle is maximized at its center. The maximum load can be decreased by redirecting routes that pass near the center to regions of lower traffic. This motivates the examination of one-turn rectilinear routing policies which we now define.

A *monotonic rectilinear routing policy* assigns to every pair (u, v) a route consisting of a monotonic rectilinear path from u to v , *i.e.*, a path comprised of a series of axis-parallel line segments such that any axis-parallel line intersects the path at most once. A *one-turn rectilinear routing policy* assigns to every pair (u, v) a monotonic rectilinear path consisting of one horizontal line segment and one vertical line segment joining u to v via an intermediate point w . Point w may coincide with u or v .

For any monotonic rectilinear routing policy P ,

$$\text{length}(\text{route}_P(p, q)) = |p_x - q_x| + |p_y - q_y|. \quad (9)$$

In general, there are two possible one-turn rectilinear routes from a given origin (u_x, u_y) to a given destination (v_x, v_y) . We refer to these as *row-first* and *column-first*, where the row-first route passes through the intermediate point (v_x, u_y) and the column-first route passes through the intermediate point (u_x, v_y) .

3 Straight-Line Routing on a Square

In this section we examine the load of straight-line routing on the unit square. These values serve as milestones against which the optimality of all other routing policies on the unit square are compared.

3.1 Average Load

By Proposition 1, the average load in the unit square under straight-line routing is equal to the expected distance between two points selected at random in the square. This value is a box integral with the following solution [17]:

$$\text{length}_{\text{avg}}(S) = \frac{2 + \sqrt{2} + 5 \ln(1 + \sqrt{2})}{15} \approx 0.5214. \quad (10)$$

By (6), the average load (and maximum load) of any routing policy on the unit square is bounded from below by (10).

3.2 Load at an Arbitrary Point

Since straight-line routing is symmetric in the x - and y -dimensions, we derive the load at an arbitrary point p located in an octant of the unit square. The load at an arbitrary point in the unit square is then easily found using the appropriate coordinate transformation.

Theorem 1. *Given a point $p = (p_x, p_y)$ such that $1/2 \leq p_y \leq p_x \leq 1$, the load at p using straight-line routing is given by*

$$\begin{aligned} \lambda_S(p) = & (1 - p_x)p_x \left[\int_{\theta=0}^{\alpha} g_1(\theta) \right] + (1 - p_x)^2 p_y \left[\int_{\theta=\alpha}^{\pi/2-\beta} g_4(\theta) \right] \\ & + (1 - p_x)p_y^2 \left[\int_{\theta=\alpha}^{\pi/2-\beta} g_3(\theta) \right] + (1 - p_y)p_y \left[\int_{\theta=\pi/2-\beta}^{\pi/2} g_2(\theta) \right] \\ & + (1 - p_y)p_y \left[\int_{\theta=\pi/2}^{\pi/2+\gamma} g_2(\theta) \right] - (1 - p_y)^2(1 - p_x) \left[\int_{\theta=\pi/2+\gamma}^{\pi-\delta} g_3(\theta) \right] \\ & + (1 - p_y)(1 - p_x)^2 \left[\int_{\theta=\pi/2+\gamma}^{\pi-\delta} g_4(\theta) \right] - p_x(1 - p_x) \left[\int_{\theta=\pi-\delta}^{\pi} g_1(\theta) \right], \quad (11) \end{aligned}$$

where expressions for α , β , γ , δ , and g_1 through g_4 are omitted for lack of space.

Expression (11) has a closed-form polylogarithmic representation (free of any trigonometric terms). The complete expression is not reproduced here due to the large number of terms but can be easily reconstructed from (11).

3.3 Maximum Load

We now derive the maximum load for straight-line routing on the unit square and show that this value is realized at the center of the square.

Theorem 2. *The maximum load for straight-line routing on the unit square is*

$$\lambda_{\max}(S) = \frac{1}{\sqrt{2}} + \frac{3}{8} \ln(\sqrt{2} + 1) - \frac{1}{8} \ln(\sqrt{2} - 1) \approx 1.1478, \quad (12)$$

realized uniquely at the center of the square.

4 One-Turn Rectilinear Routing on a Square

In this section we consider various one-turn rectilinear routing policies on the unit square and compare these against straight-line routing. Our objective in designing these policies was to reduce the maximum load by redirecting routes for particular regions of origin-destination pairs away from high-traffic areas and towards low-traffic areas while maintaining a low stretch factor.

4.1 Average Load

Theorem 3. *The average load for any monotonic rectilinear routing policy on the unit square is $2/3$.*

Proof. By Proposition 1 and (9), the average load is equal to the average ℓ_1 distance between two points in the unit square. This value is

$$\lambda_{\text{avg}}(P) = \int_0^1 \int_0^1 \int_0^1 \int_0^1 |u_x - v_x| + |u_y - v_y| dv_y dv_x du_y du_x = \frac{2}{3}. \quad \square \quad (13)$$

4.2 Average Stretch Factor

It is straightforward to see that the maximum stretch factor for any monotonic rectilinear routing policy is $\sqrt{2}$. We now consider the average stretch factor.

Theorem 4. *The average stretch factor for any monotonic rectilinear routing policy P on the unit square is*

$$\text{str}_{\text{avg}}(P) = \frac{1}{6} \left(10 \ln(2 + \sqrt{2}) + 2\sqrt{2} - 4 - 5 \ln(2) \right) \approx 1.2737. \quad (14)$$

4.3 Diagonal Rectilinear Routing

We define a routing policy in terms of the partition of the unit square induced by its two diagonals. Let R_1 through R_4 denote the four regions of the partition such that R_1 is at the bottom of the square and the regions are numbered in clockwise order. If the origin lies in R_1 or R_3 , the row-first route is selected. Otherwise, the column-first route is selected. We refer to this routing policy, denoted P_D , as *diagonal rectilinear routing*.

As we did in Section 3.2, we derive the load at an arbitrary point p located in an octant of the unit square since P_D is symmetric in the x - and y -dimensions. The load at an arbitrary point in the unit square is then easily found using the appropriate coordinate transformation.

Theorem 5. *Given a point $p = (p_x, p_y)$ such that $0 \leq p_y \leq p_x \leq 1/2$, the load at p using diagonal rectilinear routing is*

$$\lambda_{P_D}(p) = 2p_x^3 - 5p_x^2 + \frac{7}{2}p_x - 2p_x p_y + \frac{3}{2}p_y - 3p_y^2 + 2p_y^3. \quad (15)$$

Proof. Let $u = (u_x, u_y)$ denote the origin and let $v = (v_x, v_y)$ denote the destination. The relative positions of u , v , and p can be divided into seven cases such that the load at p corresponds to the sum of the measure of the regions of possible origin-destination combinations in each case. In Cases 1a through 1d, u lies in region R_1 or R_3 and, consequently, the row-first route is selected. In Cases 2a through 2c, u lies in region R_2 or R_4 and, consequently, the column-first route is selected. See Fig. 1. The theorem follows by summing the contribution to load




Fig. 1. Illustration in support of Theorem 5. The white dot denotes point p .

in each of these cases. Details are omitted for lack of space. \square

It can be shown that (15) is maximized when $p_x = \frac{5}{6} - \frac{1}{3}\sqrt{3 - \frac{1}{2}\sqrt{11}} \approx 0.4472$ and $p_y = \frac{2}{3} - \frac{\sqrt{11}}{6} \approx 0.1139$, with load $\lambda_{\max}(P_D) = \frac{1}{27}[\sqrt{11} - \frac{31}{2}] \approx 0.7771$.

4.4 Additional Policies Considered

We describe additional one-turn rectilinear routing policies considered. In each case, the maximum load was shown to be strictly greater than that of diagonal rectilinear routing. Recall that all one-turn rectilinear routing policies have equal average load (Theorem 3).

Equal Distribution A simple initial strategy to consider is to assign to each origin-destination pair (u, v) the row-first route. For any point $p \in [0, 1]^2$, $\lambda_{P_R}(p) = \lambda_P(p)$, where P_R denotes the row-first routing policy and P denotes any policy that assigns the pairs (u, v) and (v, u) different one-turn rectilinear routes for all u and v .

Outer Turn Consider the routing policy that selects the one-turn rectilinear route whose intermediate point is furthest from the center of the square. If the two intermediate points are equidistant from the origin, then a route is assigned as in the equal distribution policy.

Grid-Based Regions Divide the unit square into nine rectangular regions whose boundaries intersect the x - and y -axes at 0 , k , $1 - k$, and 1 , respectively, for some fixed $k \in [0, 1/2]$. There are three types of regions: corner, mid-boundary,

and one central region. If the origin is located in a mid-boundary region and the destination is in a non-adjacent corner region then select the one-turn rectilinear route that avoids passing through the central region. Similarly, a route from a corner regions to a non-adjacent mid-boundary region must avoid passing through the central region. For all other origin-destination pairs, routes are assigned as in the equal distribution policy. A second grid-based routing policy is defined by adding the constraint that a route from a mid-boundary region to an adjacent mid-boundary region must also avoid passing through the central region.

Line Division Let l denote the line passing through the origin u and destination v . If l does not pass through the center of the square, c , select the one-turn rectilinear route whose intermediate point is opposite l from c . If l passes through c , then a route is assigned as in the equal distribution policy.

Random Intermediate Point A random intermediate point connected to the origin and destination by straight-line routes results in load exactly twice that of straight-line routing. Perhaps a better strategy is that described by Busch et al. [3] which can be generalized to convex regions; an intermediate point is selected at random on the perpendicular bisector of the origin and destination. Note that although this policy involves a single turn, it is not a rectilinear routing policy.

Summary Table 1 summarizes bounds on average and maximum load for each of the above routing policies and straight-line routing. We derived the load at an arbitrary point in the unit square for five of these policies; the corresponding plots are illustrated in Fig. 2. The diagonal rectilinear routing policy, P_D , achieves the lowest maximum load, significantly lower than the maximum load of straight-line routing and not much greater than the lower bound.

routing policy	λ_{avg}	λ_{max}	routing policy	λ_{avg}	λ_{max}
straight-line S	0.5214	1.1478	grid-based (1)	2/3	7/8
diagonal P_D	2/3	0.7771	grid-based (2)	2/3	0.8541
equal dist. P_R	2/3	1	line division	2/3	$\geq 7/8$
outer turn	2/3	≥ 0.8977	lower bound		0.7076

Table 1. Comparing routing policies on the unit square

5 Lower Bounds on Load for One-Turn Rectilinear Routing Policies

Naturally, no monotonic rectilinear routing policy can have a maximum load less than the average load of $2/3$. In this section we establish a stronger lower bound on the maximum load of any one-turn rectilinear routing policy.

Theorem 6. *No one-turn rectilinear routing policy can guarantee a maximum load less than 0.7076.*




Fig. 2. These plots display $\lambda_P(p)$ for $p \in [0, 1]^2$ for five routing policies: (left to right) straight-line S , equal distribution P_R , grid-based (two policies), and diagonal P_D .

6 Optimal Randomized One-Turn Rectilinear Routing Policies

In this section we give a characterization of the optimal randomized one-turn rectilinear strategy as the solution of an optimization problem and provide an efficient procedure for approximating it. A deterministic one-turn rectilinear strategy is equivalent to a function $P : [0, 1]^4 \rightarrow \{0, 1\}$ where $P(u, v, s, t) = 1$ iff the route from (u, v) to (s, t) uses the column-first path. (If $u = s$ or $v = t$, we define $P(u, v, u, t) = P(u, v, s, v) = 1$.) We can generalize this to randomized rectilinear schemes by considering $Q : [0, 1]^4 \rightarrow [0, 1]$ where if $Q(u, v, s, t) = q$ then a packet travelling from point (u, v) to (s, t) takes the the column-first path with probability q and the row-first path with probability $1 - q$. A formula for the expected $\lambda(x, y)$ at a point (x, y) can be found easily.

The optimal strategy is given by the solution to the following optimization problem: $\min_Q \max_{(x,y)} \lambda(x, y)$. While we can't directly solve this problem we can approximate it by considering finer and finer partitions of the square into n^2 $1/n$ by $1/n$ subsquares and giving a strategy for all packets routing between each pair of subsquares. Now our problem is equivalent to finding a randomized one-turn rectilinear routing strategy for an $n \times n$ grid that minimizes the number of packets using any particular node of the grid under an all-to-all communication pattern.

Let p_{ijkl} , $1 \leq i, j, k, l \leq n$, be the probability that a packet starting in subsquare (i, j) going to subsquare (k, l) uses the column-first path and let the maximum expected load at any point in subsquare (r, s) be $\lambda(r, s)$.

An upper bound on $\lambda(r, s)$ is easily derived and our problem now reduces to $\min_{p_{ijkl}} \max_{r,s} \lambda(r, s)$, which is equivalent to the following linear program with $n^4 + 1$ variables and $2n^4 + n^2$ constraints (solvable in polynomial time):

$$\begin{aligned}
 & \text{Minimize } z \\
 & \text{Subject to} \\
 & 0 \leq p_{ijkl} \leq 1, \quad 1 \leq i, j, k, l \leq n, \\
 & z - \lambda(r, s) \geq 0, \quad 1 \leq r, s \leq n.
 \end{aligned}$$

Table 2 shows an upper bound on the maximum load achieved by the strategy obtained by using an $n \times n$ grid to approximate the unit square for $2 \leq n \leq 12$. The results indicate that the optimal strategy achieves a maximum load

n	2	3	4	5	6	7	8	9	10	11	12
max. load	1.0000	0.8889	0.8264	0.8009	0.7813	0.7759	0.7650	0.7610	0.7530	0.7499	0.7446

Table 2. Approximations to the optimal randomized strategy using $n \times n$ grids.

of approximately 0.74. The solutions were found using the CVXOPT convex optimization package [5]. We were unable to obtain results for larger n due to memory limitations.

References

1. Bose, P., Morin, P., Stojmenovic, I., Urrutia, J.: Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Networks* **7** (2001) 609–616
2. Broch, J., Johnson, D., Maltz, D.: The dynamic source routing protocol for mobile ad hoc networks (1998) Internet-draft, draft-ietf-manet-dsr-00.txt.
3. Busch, C., Magdon-Ismael, M., Xi, J.: Oblivious routing on geometric networks. In: *Proc. ACM SPAA*. Volume 17. (2005)
4. Chung, F.R.K., Jr., E.G.C., Reiman, M.I., Simon, B.: The forwarding index of communication networks. *IEEE Trans. Inf. Th.* **33** (1987) 224–232
5. Dahl, J.: Cvxopt: A python package for convex optimization. In: *Proc. Eur. Conf. Op. Res.* (2006)
6. Gao, J., Zhang, L.: Load balanced short path routing in wireless networks. In: *Proc. IEEE INFOCOM*. Volume 23. (2004) 1099–1108
7. Gao, J., Zhang, L.: Tradeoffs between stretch factor and load balancing ratio in routing on growth restricted graphs. In: *Proc. ACM PODC*. (2004) 189–196
8. Gupta, P., Kumar, P.R.: The capacity of wireless networks. *IEEE Trans. Inf. Th.* **46** (2000) 388–404
9. Heydemann, M.C., Meyer, J.C., Sotteau, D.: On forwarding indices of networks. *Disc. App. Math.* **23** (1989) 103–123
10. Hyttiä, E., Lassila, P., Virtamo, J.: Spatial node distribution of the random waypoint mobility model with applications. *IEEE Trans. Mob. Comp.* **6** (2006) 680–694
11. Jinyang, L., Blake, C., Couto, D.D., Lee, H., Morris, R.: Capacity of ad hoc wireless networks. In: *Proc. ACM MOBICOM*. (2001)
12. Kranakis, E., Singh, H., Urrutia, J.: Compass routing on geometric networks. In: *Proc. CCCG*. (1999) 51–54
13. Meyer auf der Heide, F., Schindelbauer, C., Volbert, K., Grunwald, M.: Energy, congestion, and dilation in radio networks. In: *Proc. ACM SPAA*. (2002) 230–237
14. Park, V., Corson, S.: A highly adaptive distributed routing algorithm for mobile wireless networks. In: *Proc. IEEE INFOCOM*. (1997) 1405–1413
15. Pham, P.P., Perreau, S.: Performance analysis of reactive shortest path and multi-path routing mechanism with load balance. In: *Proc. IEEE INFOCOM*. (2003) 251–259
16. Popa, L., Rostami, A., Karp, R.M., Papadimitriou, C., Stoica, I.: Balancing traffic load in wireless networks with curveball routing. In: *Proc. ACM MOBIHOC*. (2007)
17. Santaló, L.A.: *Integral Geometry and Geometric Probability*. Cambridge University Press (2004)
18. Stojmenovic, I.: Position based routing in ad hoc networks. *IEEE Comm. Mag.* **40** (2002) 128–134