

Memoryless Routing in Convex Subdivisions: Random Walks are Optimal

Dan Chen^a, Luc Devroye^b, Vida Dujmović^a, Pat Morin^a

^a*School of Computer Science, Carleton University, 1125 Colonel By Drive
Ottawa, Ontario, Canada, K1S 5B6*

^b*School of Computer Science, McGill University, 3480 University Street
Montreal, Québec, Canada, H3A 2A7*

Abstract

A *memoryless* routing algorithm is one in which the decision about the next edge on the route to a vertex t for a packet currently located at vertex v is made based only on the coordinates of v , t , and the neighbourhood, $N(v)$, of v . The current paper explores the limitations of such algorithms by showing that, for any (randomized) memoryless routing algorithm \mathcal{A} , there exists a convex subdivision on which \mathcal{A} takes $\Omega(n^2)$ expected time to route a message between some pair of vertices. Since this lower bound is matched by a random walk, this result implies that the geometric information available in convex subdivisions does not reduce the worst-case routing time for this class of routing algorithms. The current paper also shows the existence of triangulations for which the RANDOM-COMPASS algorithm proposed by Bose *et al* (2002, 2004) requires $2^{\Omega(n)}$ time to route between some pair of vertices.

1. Introduction

In recent years, motivated primarily by the proliferation of wireless networks and GPS devices, much research has been done on routing algorithms for geometric networks [3]. In this research a network is modelled as a geometric graph $G = (V, E)$ whose vertex set V is a set of points in \mathbb{R}^2 . We say that a routing algorithm \mathcal{A} *works* for G if, for any pair of vertices $s, t \in V$, the algorithm always finds a walk from s to t in a finite number of steps.

The research on geometric routing algorithms largely focuses on utilizing geometric properties of a class of geometric graphs to reduce the complexity of, and information required by, routing algorithms. For example, when G is a planar graph, then an algorithm of Kranakis *et al* [4] works for G and requires no preprocessing of G or additional state information at the vertices of G and requires only a packet header that can store the locations of a constant number of vertices. Routing algorithms like this, that require only storing $O(\log n)$ bits and $O(1)$ vertex identifiers are sometimes called *constant-memory routing algorithms*.

An extremely general result, but that does not rely on geometry, based on logspace construction of universal exploration sequences, shows that, using a header containing only $O(\log n)$ bits, one can visit all the vertices of any graph (and hence reach t) in a polynomial number of steps [9]. However, this algorithm is far too slow for practical use.

A particularly interesting and restricted class of routing algorithms are so-called memoryless routing algorithms. For the purposes of this paper, a *memoryless* routing algorithm is one in which the decision about the next edge on the route to t for a packet currently located at node v is based only on the coordinates of v , t , and the neighbourhood, $N(v)$, of v . More precisely, a deterministic memoryless routing algorithm is a function $f : \mathbb{R}^2 \times \mathbb{R}^2 \times (\mathbb{R}^2)^+ \rightarrow \mathbb{R}^2$ that satisfies $f(v, t, N(v)) \in N(v)$.

We note that one can find different definitions of the term “memoryless routing algorithm” in other papers, some of which do not even allow the routing algorithm to know the destination t ; these are perhaps better called memoryless traversal algorithms since, in order to work, they must visit every vertex of G . The “memoryless” restriction in the current paper is not quite as strict. It allows the algorithm a read-only memory that stores the location of t .

Note that a memoryless routing algorithm makes each routing step without using information obtained in previous routing steps and without any global information about G . Memoryless algorithms are different from *oblivious* routing algorithm [7, Section 4.2] which select a path from s to t having total knowledge of G (but without knowledge of other source/destination pairs).

Bose and Morin [2] show that if G is Delaunay a triangulation¹ or a regular triangulation then deterministic memoryless routing algorithms, named GREEDY and COMPASS, respectively, work for G . Bose *et al* [1] subsequently show a stronger result; a deterministic memoryless routing algorithm, named GREEDY-COMPASS, works for any triangulation G .

Memoryless routing algorithms are so simple, elegant, and practical that researchers have spent considerable effort designing geometric embeddings of graphs so that memoryless routing algorithms can be applied to the resulting embeddings. A famous example in this vein is due to Leighton and Moitra [5] who prove that every 3-connected planar (non-geometric) graph \tilde{G} has an embedding G as a geometric graph with vertices in \mathbb{R}^2 such that GREEDY works on G . Specifically, there is one embedding, G , such that, for every source destination pair s, t , the GREEDY algorithm always succeeds in routing from s to t . The combination of the embedding and routing algorithm represents a form of *compact routing* [10].

Unfortunately, deterministic memoryless routing algorithms have severe limitations. These stem from the fact that these algorithms can not visit the same vertex more than once without looping forever. Bose *et al* [2, Theorem 2] show

¹A *triangulation* is a plane geometric graph all of whose faces, except the outer face, are triangles, and whose outer face is the complement of a triangle. Delaunay triangulations and regular triangulations are special types of triangulations. For details, consult [8].

that there exists 17 convex subdivisions², G_1, \dots, G_{17} , each with 17 vertices such that any deterministic memoryless routing algorithm does not work for at least one of these subdivisions. Thus, convex subdivisions form a class of geometric graphs that are too rich for deterministic memoryless routing algorithms [1].

The same authors [1, 2] observe that randomization can be used to overcome this limitation. A *randomized memoryless* routing algorithm is one in which the decision about the next edge on the route to t for a packet currently located at node v is based only on v , t , the neighbourhood, $N(v)$, of v , and a sequence B of fresh random bits. More precisely, a randomized memoryless routing algorithm is defined by a function $f : \mathbb{R}^2 \times \mathbb{R}^2 \times (\mathbb{R}^2)^+ \times \{0, 1\}^\infty \rightarrow \mathbb{R}^2$ that satisfies $f(v, t, N(v), B) \in N(v)$ for all inputs. The final argument B is a sequence of random bits that are chosen fresh for each step taken by the routing algorithm. Bose *et al* describe a randomized memoryless algorithm, named RANDOM-COMPASS, that uses one random bit per step and works for any convex subdivision. They do not analyze the efficiency of RANDOM-COMPASS except to note that, for some convex subdivisions G , and some pairs $s, t \in V$, the expected number of steps taken by RANDOM-COMPASS when routing from s to t is $\Omega(|V|^2)$.

Observe that, by the theory of random walks (c.f. [7, Theorem 6.6]), the expected time required for a random walk on G to travel from a particular vertex s to a particular vertex t is $O(n^2)$. Therefore, a random walk is at least as efficient, in the worst case, as the RANDOM-COMPASS algorithm. Nevertheless, one might expect that RANDOM-COMPASS is more likely to find short routes, since it uses geometry to find a route that is specifically directed towards the target vertex t . Thus, we might intuit that RANDOM-COMPASS is a heuristic that is usually better than a random walk and never much worse.

In the current paper, we show that this intuition about RANDOM-COMPASS could not be further from the truth. Indeed, for any $n > 0$, there exists a convex subdivision (in fact, a triangulation) G with n vertices and having two vertices s and t such that the expected number of steps taken by RANDOM-COMPASS when routing from s to t is $2^{\Omega(n)}$. This triangulation has diameter 3.

Next we study whether *any* randomized memoryless routing algorithm for convex subdivisions can outperform a random walk. We show that, for any randomized memoryless routing algorithm \mathcal{A} and any n , there exists a convex subdivision $G = G(\mathcal{A}) = (V, E)$ of size n and a pair of vertices $s, t \in V$ such that the expected number of steps taken by \mathcal{A} when routing from s to t is $\Omega(n^2)$. Therefore, \mathcal{A} 's performance is, in some cases, no better than that of a random walk.

²A *convex subdivision* is a plane geometric graph all of whose faces, except the outer face, are convex polygons, and whose outer face is the complement of a convex polygon.




Figure 1: The RANDOM-COMPASS algorithm chooses the next vertex at random among $ccw_t(v)$ and $cw_t(v)$.

2. A Bad Example for Random-Compass

The RANDOM-COMPASS algorithm works by using a coin toss to select among the (at most two) neighbours $ccw_t(v)$ and $cw_t(v)$ of the current node v that make the minimum counterclockwise and clockwise angle, respectively, with the segment vt (see Figure 1.a). When applied on a convex subdivision $G = (V, E)$, Bose *et al* show that, in the directed graph G' that contains the edges $(v, cw_t(v))$ and $(v, ccw_t(v))$ for all $v \in V$, there exists at least one directed path $P(v, t)$ from every vertex v to t (see Figure 1.b). This, and Wald's Equation, immediately imply that the expected time to reach t from any vertex is at most 2^n ; from any vertex v , RANDOM-COMPASS has probability at least $1/2^{|P(v,t)|} \geq 1/2^{n-1}$ of reaching t by following $P(v, t)$, and the expected number of steps it takes on $P(v, t)$ before falling off $P(v, t)$ is at most 2.

The example in Figure 2 shows that the above analysis of RANDOM-COMPASS, although very coarse, is about the best one can do. It shows a geometric graph G whose vertex set has size $n = 4k + 1$ and whose vertices are organized as a central vertex t and four paths leading from the outer face to t . The space between these paths is triangulated so that, at any point, RANDOM-COMPASS chooses between an edge that leads one step closer to t or that returns to the outer face.

If we consider the directed graph G' defined above, then we see that, at any point if the packet is at some distance $i > 1$ from t , it can, with equal probability, move to a vertex of distance $i - 1$ or move to a vertex (on the outer face) of distance k . If we denote by T_i the expected number of steps required by RANDOM-COMPASS to reach t given that it is currently at distance i from t , we see that

$$T_i = \begin{cases} 1 & \text{for } i = 1 \\ 1 + (1/2)T_{i-1} + (1/2)T_k & \text{for } i \in \{2, \dots, k\} \end{cases}$$




Figure 2: A graph in which RANDOM-COMPASS has expected running time $\Omega(2^{n/4})$.

Expanding the value of T_k gives

$$\begin{aligned}
 T_k &= 1 + (1/2)T_k + (1/2)T_{k-1} \\
 &= 1 + (1/2)T_k + 1/2 + (1/4)T_k + (1/4)T_{k-2} \\
 &= 1 + (1/2)T_k + 1/2 + (1/4)T_k + (1/4) \\
 &\quad + \dots + (1/2^{k-2}) + (1/2^{k-1})T_k + (1/2^{k-1})T_1 \\
 &= 2 - 1/2^{k-1} + (1 - 1/2^{k-1})T_k,
 \end{aligned}$$

and rewriting this gives $T_k = 2^k - 1 = \Omega(2^{n/4})$. This proves:

Theorem 1. *For any $n > 1$, there exists a triangulation G having two vertices s and t such that the expected number of steps taken by RANDOM-COMPASS when routing from s to t is $2^{\Omega(n)}$.*

Note that the base in the exponent can be improved by using a construction with 3 paths instead of 4. In this case, the lower bound becomes $\Omega(2^{n/3})$. Furthermore, up to a factor of 2, the lower bound on Theorem 1 holds for all choices of the source vertex s since, for any vertex $s \neq t$, the expected time to route from s to t is at least $(1/2)T_k$.

3. A Lower Bound for Any Algorithm

In this section we develop an $\Omega(n^2)$ lower bound on the worst-case performance of any randomized memoryless routing algorithm \mathcal{A} . In particular, we show that for any such \mathcal{A} , there exists a geometric graph $G = G(\mathcal{A})$ having two vertices s and t such that the expected number of steps taken by \mathcal{A} when routing from s to t is $\Omega(n^2)$.

The outline of the lower bound is as follows: We start with a lemma about Markov chains whose transition graphs are paths. We show that, when starting at the midpoint of the path, there is at most one endpoint of the path that can be reached in subquadratic expected time. This lemma is relevant since, if \mathcal{A} finds itself in the interior of a path of degree 2 vertices in G , it will behave like such a Markov chain until it reaches one of the endpoints of this path.

Next, we observe how \mathcal{A} behaves on certain paths of degree 2 vertices and show that, because \mathcal{A} can only reach one endpoint of any path in subquadratic time, we can always find a subset of these paths that can be pieced together to form a convex subdivision in which \mathcal{A} takes at least quadratic expected time to route from some vertex s to some vertex t .

3.1. Markov Chains

Consider a Markov chain on $\{1, \dots, n\}$, $n > 1$, where transitions only take place between neighbors. If $p_{i,j}$ is the probability of a transition from i to j , then we have

$$p_{1,2} = p_{n,n-1} = 1,$$

$$p_{i,i+1} = 1 - p_{i,i-1} = \pi_i, 2 \leq i \leq n-1,$$

where π_2, \dots, π_{n-1} are fixed probabilities. The vector of these probabilities is denoted by π . We will set $\pi_1 = 1$, $\pi_n = 0$, to be consistent, as the extreme states are reflecting. When $\pi_i = 1/2$ for $2 \leq i \leq n-1$, we obtain a standard random walk on a finite interval with reflecting barriers.

We denote the Markov chain by $X_0, X_1, \dots, X_t, \dots$, and denote the hitting times by $T_{i,j}$:

$$T_{i,j} = \min\{t > 0 : X_t = j | X_0 = i\}.$$

For a standard random walk, it is known that

$$\mathbb{E}\{T_{i,j}\} = (j-i)^2, j \neq i, 1 \leq j, i \leq n$$

[6]. The standard random walk is in fact the best possible chain in the following sense:

Lemma 1. *For any vector of probabilities π , and any $n > 1$,*

$$\mathbb{E}\{T_{1,n} + T_{n,1}\} \geq 2(n-1)^2.$$

Proof. The lemma is obviously true if any π_i , $2 \leq i \leq n-1$, is either zero or one as that would imply that at least one of the hitting times is infinite. Thus, we assume that all probabilities are strictly in $(0, 1)$. It is also trivial if $n = 2$, so assume $n > 2$. Define

$$P_i = \frac{1}{\pi_i} - 1, Q_i = \frac{1}{1 - \pi_i} - 1,$$

and note that $P_i Q_i = 1$. If needed, we formally set $P_1 = Q_n = 0$.

We need an explicit formula for $E\{T_{1,1}\}$. Let us introduce the chains on $\{i, \dots, n\}$ with reflecting barriers at i and n , but with the same π_j values associated with non-terminal states. Let $T_{i,j}^+$ with $j \geq i$, denote the hitting time from i to j in the chain $\{1, \dots, n\}$ defined this way. Clearly,

$$T_{n-1,n-1}^+ = 2.$$

Next,

$$T_{n-2,n-2}^+ = 2 + \sum_{j \leq Z} W_j,$$

where W_j are independent lengths excursions from $n-1$ to $n-1$ on the chain $\{n-1, n\}$, and Z (possibly zero) is the number of such excursions. Obviously, Z is geometrically distributed, and $E\{Z\} = Q_{n-1}$. Because $E\{W_1\} = E\{T_{n-1,n-1}^+\} = 2$, and because Z is a stopping time, we have, by Wald's identity,

$$E\{T_{n-2,n-2}^+\} = 2 + 2Q_{n-1}.$$

This argument is easily extended by induction, and we obtain for $1 \leq i < n-1$,

$$E\{T_{i,i}^+\} = 2 + Q_{i+1}E\{T_{i+1,i+1}^+\} = 2(1 + Q_{i+1} + Q_{i+1}Q_{i+2} + \dots + Q_{i+1} \dots Q_{n-1}).$$

By flipping sides, and denoting by $T_{i,j}^-$ the hitting times from i to j for the Markov chains on $\{1, \dots, i\}$ with reflecting barriers at 1 and i , we obtain in a similar fashion, for $2 < i \leq n$,

$$E\{T_{i,i}^-\} = 2 + P_{i-1}E\{T_{i-1,i-1}^-\} = 2(1 + P_{i-1} + P_{i-1}P_{i-2} + \dots + P_{i-1} \dots P_2).$$

Furthermore, $E\{T_{2,2}^-\} = 2$.

With these calculations out of the way, we note that

$$E\{T_{i,i+1}\} = 1 + P_i \times E\{T_{i,i}^-\},$$

and

$$E\{T_{i,i-1}\} = 1 + Q_i \times E\{T_{i,i}^+\}.$$

Clearly,

$$\begin{aligned}
& \mathbb{E}\{T_{1,n} + T_{n,1}\} \\
&= \sum_{i=1}^{n-1} \mathbb{E}\{T_{i,i+1}\} + \sum_{i=2}^n \mathbb{E}\{T_{i,i-1}\} \\
&= 2(n-1) + 2 \sum_{i=2}^{n-1} (P_i + P_i P_{i-1} + \cdots + P_i \cdots P_2) \\
&\quad + 2 \sum_{i=2}^n (Q_i + Q_i Q_{i+1} + \cdots + Q_i \cdots Q_{n-1}) \\
&= 2(n-1) + 2 \sum_{i=2}^{n-1} \sum_{j=2}^i \prod_{k=j}^i P_k + 2 \sum_{i=2}^n \sum_{j=i}^{n-1} \prod_{k=i}^j Q_k \\
&= 2(n-1) + 2 \sum_{i=2}^{n-1} \sum_{j=2}^i \prod_{k=j}^i P_k + 2 \sum_{j=2}^n \sum_{i=j}^{n-1} \prod_{k=j}^i Q_k \\
&= 2(n-1) + 2 \sum_{i=2}^{n-1} \sum_{j=2}^i \prod_{k=j}^i P_k + 2 \sum_{i=2}^{n-1} \sum_{j=2}^i \prod_{k=j}^i Q_k \\
&= 2(n-1) + 2 \sum_{i=2}^{n-1} \sum_{j=2}^i \left(\prod_{k=j}^i P_k + \prod_{k=j}^i Q_k \right) \\
&\geq 2(n-1) + 4 \sum_{i=2}^{n-1} \sum_{j=2}^i \sqrt{\prod_{k=j}^i P_k \times \prod_{k=j}^i Q_k} \\
&\quad \text{(by the arithmetic-geometric mean inequality)} \\
&= 2(n-1) + 4 \sum_{i=2}^{n-1} \sum_{j=2}^i 1 \\
&\quad \text{(since } P_i Q_i = 1 \text{ for all } i \text{ in our range)} \\
&= 2(n-1) + 4 \sum_{i=1}^{n-1} (i-1) \\
&= 2(n-1) + 2n(n-1) - 4(n-1) \\
&= 2(n-1)^2,
\end{aligned}$$

which concludes the proof. \square

Next we present a simple corollary of Lemma 1 that is used in our lower bound.

Corollary 1. *Consider a random walk with reflecting barriers on $\{-n, \dots, n\}$, $n > 0$. In this chain,*

$$\max(\mathbb{E}\{T_{0,n}\}, \mathbb{E}\{T_{0,-n}\}) \geq \frac{2}{3} n^2.$$

Proof. We prove this by contradiction. Set $c = 2/3$. Assume that

$$\max(\mathbb{E}\{T_{0,n}\}, \mathbb{E}\{T_{0,-n}\}) < cn^2.$$

By Theorem 1,

$$\mathbb{E}\{T_{0,n}\} + \mathbb{E}\{T_{n,0}\} \geq 2n^2,$$

and

$$\mathbb{E}\{T_{0,-n}\} + \mathbb{E}\{T_{-n,0}\} \geq 2n^2.$$

Observe for this that 0 is not a reflecting barrier, but this makes $\mathbb{E}\{T_{0,n}\}$ only larger, so Theorem 1 does indeed apply. By our assumption, we thus have

$$\min(\mathbb{E}\{T_{n,0}\}, \mathbb{E}\{T_{-n,0}\}) > (2-c)n^2.$$

Let T be the cover time, i.e., the time to visit all states starting from state 0. It is easy to see that

$$T_{0,n} + T_{0,-n} > T = \max(T_{0,n}, T_{0,-n}) = T_{0,S} + T_{S,-S},$$

where $S \in \{n, -n\}$ is the first of the two end states reached by the Markov chain. If we condition on the history up to $T_{0,S}$, we see that

$$\mathbb{E}\{T_{S,-S}\} \geq \min(\mathbb{E}\{T_{n,0}\}, \mathbb{E}\{T_{-n,0}\}) > (2-c)n^2.$$

Thus,

$$\begin{aligned} \max(\mathbb{E}\{T_{0,n}\}, \mathbb{E}\{T_{0,-n}\}) &\geq \frac{1}{2} (\mathbb{E}\{T_{0,n}\} + \mathbb{E}\{T_{0,-n}\}) \\ &\geq \frac{1}{2} \mathbb{E}\{\max(T_{0,n}, T_{0,-n})\} \\ &= \frac{1}{2} \mathbb{E}\{T\} > (1-c/2)n^2, \end{aligned}$$

which contradicts our assumption. \square

3.2. The Lower Bound

Let \mathcal{A} be a randomized memoryless routing algorithm. Let k be an even integer, let t be the origin, and let $A = a_1, \dots, a_k$ be a path of k collinear vertices such that a_k is closer to t than any of a_1, \dots, a_{k-1} and the three points a_1, a_k, t make a left turn with $\angle a_1 a_k t$ greater than 150° degrees but less than 180° (see Figure 3.a). Let $B = b_1, \dots, b_k$ be the reflection of A through the line parallel to A that contains t (see Figure 3.b). Let $A(\alpha)$, respectively, $B(\alpha)$, denote the path A , respectively, B , rotated by an angle of α about the origin, t .

Define the *color* of a path $A' = A'(\alpha) = a'_1, \dots, a'_k$ as follows: Imagine running \mathcal{A} on the graph consisting of A' and the isolated vertex t , starting at $a'_{k/2}$. If \mathcal{A} takes $\Omega(k^2)$ expected time to reach a'_k then color $A(\alpha)$ *blue*, otherwise color $A(\alpha)$ *red*. Note that Corollary 1 implies that, if $A(\alpha)$ is red, then \mathcal{A} takes $\Omega(k^2)$ expected time to reach a_1 starting at $a_{k/2}$.




Figure 3: The chains A and B .




Figure 4: Two blue chains $A(\alpha)$ and $B(\alpha)$.

Intuitively, a path is red (getting hotter — closer to t) if \mathcal{A} could move quickly from $a_{k/2}$ to a_k . A path is blue (getting cooler — further from t) if \mathcal{A} could move quickly to a_1 . Define the color (red or blue) of a path $B(\alpha)$ in the same way.

Lemma 2. *If there exists α such that $A(\alpha)$ and $B(\alpha)$ are both blue, then there exists a convex subdivision $G = (V, E)$ with $|V| = 2k + 1$ with vertices $s, t \in V$ such that \mathcal{A} takes $\Omega(k^2)$ steps when routing from s to t .*

Proof. Let $A' = A(\alpha) = a'_1, \dots, a'_k$ and $B' = B(\alpha) = b'_1, \dots, b'_k$. The convex subdivision G consists of A' and B' as well as the edges $a'_1 b'_1$, $a'_k t$ and $b'_k t$ (see Figure 4). Since $A(\alpha)$ and $B(\alpha)$ both blue, applying \mathcal{A} to route from $a'_{k/2}$ to t will require $\Omega(k^2)$ expected steps. \square

Lemma 3. *If there exists α such that $A(\alpha)$ and $B(180 + \alpha)$ are both red, then there exists a convex subdivision $G = (V, E)$ with $|V| = 2k + 1$ with vertices $s, t \in V$ such that \mathcal{A} takes $\Omega(k^2)$ steps when routing from s to t .*

Proof. Let $A' = A(\alpha) = a'_1, \dots, a'_k$ and $B' = B(180 + \alpha) = b'_1, \dots, b'_k$. The convex subdivision G consists of A' and B' as well as the edges $a'_k b'_k$, $a'_1 t$ and




Figure 5: Two red chains $A(\alpha)$ and $B(180 + \alpha)$.

$b'_1 t$ (see Figure 5). Since $A(\alpha)$ and $B(180 + \alpha)$ are red, applying \mathcal{A} to route from $a'_{k/2}$ to t will require $\Omega(k^2)$ expected steps. \square

Theorem 2. *For any integer $k > 0$ and any memoryless routing algorithm \mathcal{A} , there exists a convex subdivision $G = (V, E)$ with $|V| = \Theta(k)$ having vertices $s, t \in V$ such that \mathcal{A} takes $\Omega(k^2)$ steps when routing from s to t .*

Proof. If either of Lemma 2 or Lemma 3 apply to \mathcal{A} then the proof is complete. Otherwise, observe that the exclusion of these two lemmata implies that, for any α , at least one of $A(\alpha)$ and $B(\alpha)$ is blue. To see this, note that if $A(\alpha)$ is red, then (the exclusion of) Lemma 3 implies that $B(\alpha + 180)$ is blue, so (the exclusion of) Lemma 2 implies that $A(\alpha + 180)$ is red, so (the exclusion of) Lemma 3 implies that $B(\alpha)$ is blue.

Therefore, there exists 3 blue chains $X = x_1, \dots, x_k$, $Y = y_1, \dots, y_k$, and $Z = z_1, \dots, z_k$ where $X \in \{A(0), B(0)\}$, $Y \in \{A(120), B(120)\}$ and $Z \in \{A(240), B(240)\}$. We can then take G to be the graph containing X , Y , and Z , as well as the edges $x_1 y_1$, $y_1 z_1$, $z_1 x_1$, $x_k y_k$, $y_k z_k$, $z_k x_k$, $x_k t$, $y_k t$, $z_k t$ (see Figure 6). Because X , Y , and Z are all blue, the expected number of steps required to route from $x_{k/2}$ to t using \mathcal{A} is $\Omega(k^2)$.

All that remains is to verify that G is indeed a convex subdivision. This is readily established using the fact that the angles $\angle x_1 x_k t$, $\angle y_1 y_k t$, and $\angle z_1 z_k t$, are all between 150 and 180 degrees. \square

4. Conclusions

We have shown that the RANDOM-COMPASS algorithm takes exponential expected time to route on some convex subdivisions and that, for any randomized memoryless routing algorithm \mathcal{A} there is a convex subdivision G and a pair of vertices s and t such that \mathcal{A} takes a quadratic expected number of steps when routing from s to t .

Open Problem 1. *The current upper bound for the expected time required by RANDOM-COMPASS on convex subdivisions is $O(2^n)$ and the lower bound is $\Omega(2^{n/3})$. Close this gap.*

Open Problem 2. *A random walk on G routes any message in $O(n^2)$ expected time but requires $O(\log d)$ random bits when located at a vertex of degree d . Is there a randomized memoryless routing algorithm for routing on convex subdivisions that uses $O(1)$ random bits per step and that routes any message in $O(n^2)$ expected time?*




Figure 6: Three blue chains

References

- [1] P. Bose, A. Brodnik, S. Carlsson, E. D. Demaine, R. Fleischer, A. López-Ortiz, P. Morin, and J. I. Munro. Online routing in convex subdivisions. *International Journal of Computational Geometry and Applications*, 12(4):283–296, 2002.
- [2] P. Bose and P. Morin. Online routing in triangulations. *SIAM Journal on Computing*, 33(4):937–951, 2004.
- [3] S. Giordano and I. Stojmenović. Position based routing algorithms for ad hoc networks: A taxonomy. In X. H. X. Cheng and D.-Z. Du, editors, *Ad Hoc Wireless Networking*, pages 103–136. Kluwer, 2003.
- [4] E. Kranakis, H. Singh, and J. Urrutia. Compass routing on geometric networks. In *Proceedings of the 11th Canadian Conference on Computational Geometry (CCCG’99)*, 1999.

- [5] T. Leighton and A. Moitra. Some results on greedy embeddings in metric spaces. In *FOCS*, pages 337–346. IEEE Computer Society, 2008.
- [6] J. W. Moon. Random walks on random trees. *Journal of the Australian Mathematical Society*, 15:42–53, 1973.
- [7] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
- [8] A. Okabe, B. Boots, and K. Sugihara. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Wiley, New York, 1992.
- [9] O. Reingold. Undirected st-connectivity in log-space. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC2005)*, pages 376–385, 2005.
- [10] J. van Leeuwen and R. B. Tan. Compact routing methods: A survey. In P. Flocchini, B. Mans, and N. Santoro, editors, *SIROCCO*, pages 99–110. Carleton University Press, 1994.