

Concurrent Geometric Multicasting

Jordan Adamek
Kent State University, Kent, Ohio

James Scott Robinson
Kent State University, Kent, Ohio

Mikhail Nesterenko
Kent State University, Kent, Ohio

Sébastien Tixeuil
UPMC Sorbonne Universités & IUF, Paris, France

ABSTRACT

We present *MCFR*, a multicasting concurrent face routing algorithm that uses geometric routing to deliver a message from source to multiple targets. We describe the algorithm's operation, prove it correct, estimate its performance bounds and evaluate its performance using simulation. Our estimate shows that *MCFR* is the first geometric multicast routing algorithm whose message delivery latency is independent of network size and only proportional to the distance between the source and the targets. Our simulation indicates that *MCFR* has significantly better latency and reliability than existing algorithms.

KEYWORDS

Multicasting, Geometric Routing, Wireless Algorithms.

ACM Reference format:

Jordan Adamek, Mikhail Nesterenko, James Scott Robinson, and Sébastien Tixeuil. 2018. Concurrent Geometric Multicasting. In *Proceedings of International Conference on Distributed Computing and Networking, Varanasi, India, January 2018 (ICDCN 2018)*, 10 pages. DOI: 10.1145/3154273.3154314

1 INTRODUCTION

Geometric routing. Geometric routing is transmitting a message from the sender to the targets on the basis of the geometric locations of the nodes. Geometric routing offers a number of advantages: it does not require each node to maintain routing information beyond its immediate neighborhood; it can be stateless if no information is retained at the node as it forwards the message; its message size can be kept constant as each message carries limited amount of data, independent of network size. These properties make geometric routing attractive in fast changing networks such as vehicular [17] or resource poor networks such as wireless sensor [2].

Unicasting face routing. In *unicasting*, a single source sends a message to a single target whose coordinates are known to the source. The simplest form of unicast geometric routing is greedy. In *greedy routing* [12], each intermediate node forwards the message to its neighbor that is the closest to the target. Pure greedy routing

fails if the message encounters *local minimum*: a node that does not have neighbors closer to the target. A *sequential* geometric routing algorithm, such as the classic GFG/GPSR [5, 18], routes a single message in the *greedy mode* until a local minimum is encountered. The algorithm then switches to *recovery mode*, which involves traversing the faces of a planar subgraph of the original communication graph. Specifically, the algorithm traverses the faces that intersect the line that connects the source and the target.

There are several disadvantages of sequential face traversal of GFG and similar algorithms. A sequential algorithm cannot predict which traversal direction results in shorter distance. Hence, in the worst case, the latency of message delivery of such algorithms is proportional to the network size. Furthermore, such algorithms have low reliability: a single transmission fault results in complete message delivery failure. Last, sequential algorithms are unable to determine if the target is disconnected from the source, possibly resulting in the message remaining in the network arbitrarily long. A *concurrent* unicast algorithm CFR [8] sends a pair of messages to traverse the source-target line concurrently. This naturally selects the shortest face traversal direction. Moreover, the second message provides greater robustness in case of message loss. The meeting of these messages signifies the end of face traversal and automatically determines whether the target is connected to the source. These advantages are offset by greater message cost.

Multicasting. Geometric multicasting requires the source to transmit the same message to several targets [26]. The source knows the coordinates of the targets. Unicasting to each target separately may be inefficient if several targets are located near each other since multiple redundant messages are sent to similar locations. A number of multicasting algorithms presented in the literature [7, 24, 30] optimize the routes the messages take so that a single message is routed to multiple targets as long as possible. Specifically, LGS [7] computes a minimum Euclidean length spanning tree rooted in the source and containing all the targets. A Euclidean Steiner tree is a minimum length spanning tree with virtual nodes. The total length of a Steiner tree may be half that of the minimum Euclidean length spanning tree [16]. However, its computation is NP-hard. Sophisticated routing schemes of PBM [24] and GMP [30] compute an approximation of the Euclidean Steiner tree at every intermediate node and then unicast the message along this tree. In both GMP and PBM, the message is geometric unicast towards the node, real or virtual, that roots the subtree of the target destinations. As the message approaches a virtual node, it may be advantageous to *split* the message and route separate messages towards different destination groups. GMP is shown to be more efficient than PBM in route selection and Steiner tree computation [30]. Overall, all these multicasting algorithms use sequential face routing to recover from a

*This work was supported in part by LINCS.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICDCN 2018, Varanasi, India

© 2018 ACM. 978-x-xxxx-xxxx-x/YY/MM...\$15.00

DOI: 10.1145/3154273.3154314

local minimum. Hence, they are prone to inefficient route selection, message loss and inability to detect network disconnection.

Geocasting. A problem closely related to multicasting is geocasting. In geocasting, the source needs to deliver a message to every node in a particular geographic region. The source is unaware of the individual target coordinates. There are geocasting algorithms based on sequential face traversal [5, 22, 22, 27]. There is also a concurrent geocasting algorithm [1]. Despite problem similarity, the geocasting algorithms do not know target coordinates, and to reach them all, they flood the geocast area with messages. Hence, geocast algorithms cannot be used for efficient multicasting.

Abstract and concrete simulation. In *abstract simulation*, message transmission is executed as an atomic step while physical details of radio communication are not considered. This simulation focuses on the performance of the algorithm itself. One of the major performance characteristics of abstract simulation of geometric routing algorithms is their performance at various network densities. The message delivery latency of all existing algorithms declines in mid-range density called *critical density*. Alternatively, in *concrete simulation*, many details of message transmission are modeled: physical and link level operation, radio interference, message loss, signal strength, etc. Both simulation approaches, abstract and concrete, have merits as abstract simulation focuses on the properties of algorithm itself regardless of the particular wireless technology, which may potentially change in the future. Concrete simulation studies the impact of the engineering details needed to implement the algorithm. Arguably, one of the most popular geometric routing algorithms GPSR [18] is a concrete implementation of GFG [5]. In this paper, we use both abstract and concrete simulation to evaluate the performance of multicasting algorithms.

Our contribution. We present *MC \mathcal{FR}* : the first concurrent multicasting face routing algorithm. The algorithm computes an approximation of Steiner tree of the targets and then concurrently routes the message around all the faces that intersect this Steiner tree. We prove it correct, provide asymptotic measures of its latency and message complexity and then evaluate its performance through simulation. We show that *MC \mathcal{FR}* is the first known multicasting algorithm whose latency does not depend on the network size. It only depends on the distance between the source and the targets. We evaluate the performance of *MC \mathcal{FR}* using abstract and concrete simulation. The abstract simulation shows that, unlike existing algorithms, *MC \mathcal{FR}* delivery latency is insensitive to network density. *MC \mathcal{FR}* outperforms existing algorithms at all densities; however, its advantage is particularly pronounced in critical density range where the other algorithms perform poorly. The concrete simulation indicates that the delivery ratio of *MC \mathcal{FR}* is only minimally affected by network density. Again, *MC \mathcal{FR}* is always more reliable than other algorithms and this advantage widens where existing algorithms falter, for example at higher densities or lower signal strengths.

2 NOTATION AND DEFINITIONS

Wireless network, message and node limitations. A *wireless network* $G = (N, E)$ is represented as a graph where N is a set of nodes that are devices capable of exchanging messages wirelessly, while E is a set of edges connecting the nodes if the two adjacent nodes can send messages directly. Two such nodes are *neighbors*. The communication is bi-directional and the graph is undirected. Every node has unique planar coordinates that *embed* the graph into the geometric plane. When it is clear from the context, we refer to a graph’s embedding as just graph.

Planarity, face traversal, Steiner tree. A graph embedding is *planar* if the graph edges intersect only at vertices. For short, we call this planar embedding of a graph, a *planar graph*. A *connected planar subgraph* is a subset of vertices and their induced edges such that the resultant graph is planar and connected. Finding a maximum planar subgraph of a general graph is NP-hard [23]. However, for certain graphs, the task may be solved efficiently. A graph is *unit-disk* if a pair of its vertices u and v are neighbors if and only if the Euclidean distance between them is no more than 1. Such a graph approximates a wireless network. In such a graph, a connected planar subgraph may be found by local computation at every node using Relative Neighborhood or Gabriel Graph [5, 14, 18, 28]. *Face* is a region of the plane such that any two points of the region may be connected by a continuous curve that does not intersect the edges of the graph. A planar embedding of a finite graph divides the plane into a finite set of faces. The areas of all but one of the faces are finite. The finite area faces are *internal*. The infinite face is *external*. For example, in Figure 2, the graph has three internal faces: F , G and H and the external face.

Consider node u and its neighbors v and w . Node w is *next-right* after v if it is the next neighbor of u after v clockwise; it is *next-left* after v if it is next to it counter-clockwise. For example, in Figure 2, f is next-right neighbor of i after s . Observe that if w is next-left after v , then v is next-right after w . Node u , its two neighbors v and w and the two incident edges form *angle* $\angle vuw$. An angle *intersects* a segment of a line if at least one of its edges lies on or intersects this segment. For example, $\angle sif$ intersects segment sx . Note that we limit angle intersection to the fixed-size graph edges, not the infinite half-rays of a classic geometric angle. In a planar graph, to traverse a face, messages are routed using right- or left-hand-rule. In the *right-hand-rule*, if a node receives a message, it forwards the message to the neighbor that is next-right after the sender. In the *left-hand-rule*, the message is forwarded to the next-left neighbor. For example, in Figure 2, if i receives a right-hand-rule traversal message from s , then it forwards it to f . Two messages are *mates* if they are traversing the same face in the opposite directions. A single node is able to detect mates if the sender of each message is the receiver of the other and the traversal direction of the two messages is opposite. A *Euclidean minimum Steiner tree* connects a selected set of nodes on a plane, possibly with added *virtual nodes*, by a graph of minimum total length. The problem of computing such tree is NP-hard [15]. GeoSteiner is the most successful algorithm that computes the exact solution in reasonable time [29]. Efficient polynomial approximations are also available [3, 25]. For the rest

of the paper, we refer to the Euclidean minimum Steiner tree as just Steiner tree and ignore the fact that it is being approximated.

Message and node memory constraints. To help with routing, a message carries routing information. Only *constant size messages* are allowed. This means that the message may carry only a fixed number of node coordinates and related information. This fixed number is independent of the network size. This limitation, for example, precludes a routing algorithm from requesting the message to carry its entire route. Each message always carries the immediate sender, the node transmitting this message, the immediate receiver and the node the message is being sent to. Each node stores the coordinates of its neighbors. No other information, either temporarily or permanently, can be stored by the node. This limitation precludes nodes from maintaining extensive routing tables of the network or storing state information between message transmissions.

Steps, computations, fairness, multicasting. Every node has a *send queue SQ* that collects messages to be sent. A message is transmitted by taking it from the sender's send queue, transferring it to the receiver and processing it according to the routing algorithm. In the theoretical discussion about the algorithm, we assume that this transfer and processing is done in a single atomic *step*. The *atomicity* of the step means that it may not overlap with the steps of this or any other nodes. In the concrete simulation section, this assumption is lifted. *Computation* is a sequence of atomic steps that starts in an initial state of the algorithm. A computation is *fair* if each message, in the send queue of every node, is either transmitted or removed from this queue. That is, a message may not "get stuck" in a send queue forever. We consider only fair computations. A computation with a finite number of steps is itself *finite*. A routing algorithm is *terminating* if its every computation is finite. A terminating routing algorithm never leaves messages circulating in the network indefinitely. A *multicasting routing algorithm* ensures a message is delivered from the *source* to the set of *targets*. The source knows the coordinates of the targets and these coordinates fit in a single message. For example, in Figure 2, source s may need to send a message to targets b , d and k . To aid in navigation, a multicasting algorithm may compute a Steiner tree that includes source, targets and virtual nodes x and y .

3 ALGORITHM DESCRIPTION

The pseudocode of algorithm *MCFR* is shown in Figure 1. It operates as follows. The sender computes the Steiner tree T of targets and itself. For each angle that intersects T , it sends a pair of mates with right R and left L traversal directions. Every message carries the encoding of T . For simplicity, we assume that source itself is never a target.

Once some node n receives a message, it checks if its mate is in the send queue SQ . If the mate is found, both messages are removed and further processing stops as this completes the traversal of a face. If there is no mate, n checks if it is the target and delivers the message. After the delivery, message processing continues. Specifically, the message is forwarded along the face that it traverses. A node is *juncture* if at least one of its angles intersects T . If n is a juncture node then n *splits* the message by injecting a pair of mates

```

node s
compute  $T$  /*  $T$  is Steiner tree */
foreach  $\angle asb$  that intersects  $T$  do
  add  $L(s, T, a)$  to  $SQ$ 
  add  $R(s, T, b)$  to  $SQ$ 

node n
if receive  $L(s, T, a)$  then
  if  $R(s, T, a) \in SQ$  then
    /* found mate */
    discard  $R(s, T, a)$  from  $SQ$ 
  else
    if  $n \in T$  then
      deliver  $L(s, T, a)$  to  $n$ 
      /* let  $b$  be the next left after  $a$  */
      add  $L(s, T, b)$  to  $SQ$ 
    if  $\angle anb$  intersects  $T$  then
      /* split message */
      foreach  $\angle cnd \neq \angle anb$  that intersects  $T$  do
        /* let  $d$  be the next left after  $c$  */
        add  $R(s, T, c)$  to  $SQ$ 
        add  $L(s, T, d)$  to  $SQ$ 
  if receive  $R(s, T, a)$  then
    /* handle similar to  $L(s, T, a)$  */

```

Figure 1: *MCFR* pseudocode.




Figure 2: *MCFR* operation example.

in every angle that intersects T . Observe that the source node is always a juncture.

Let us consider an example of *MCFR* operation shown in Figure 2. Source node s computes the Steiner tree T , determines that $\angle isa$ intersects T and sends a left-hand-rule L traversing message to a , while sending R to i . This injects a pair of mates into face F . When a receives a message, it forwards it to its next-left left after s neighbor b . When i receives a message, it determines that i is a juncture node and the angle that intersects T is $\angle fji$. Node i injects a pair of mates by sending an L to f and an R to j . Meanwhile, i also forwards the originally received message to f . Node f injects mates into faces H and G and forwards the original message to b .

Node b injects messages into face G and forwards the message to a . However, then b receives a message from a . Once b receives a message from a , it inspects its SQ , finds its mate and destroys both messages. This completes the traversal of face F . The operation of $MCFR$ continues until all faces that intersect T are traversed.

4 CORRECTNESS PROOF AND EFFICIENCY BOUNDS

Correctness proof. Let us introduce some notation to aid in the correctness discussion. A node is *segment-visited*, or just *visited*, with respect to a particular face if it was visited during the traversal of this face. It is *unvisited* otherwise. A *visited segment* of a face is a sequence of neighbor nodes that are visited. A *segment-border* of a visited segment, with respect to a particular face, is a visited node with a neighbor that is either (i) unvisited or (ii) carrying a message for the border node. A non-border visited node is *segment-internal*. Note that an edge in a planar graph is adjacent to two faces. Thus, a node may be visited in one face but unvisited in another. Two faces are *adjacent* if they share a common juncture node. Two faces F and G are *juncture connected* if there exists a sequence of adjacent faces that starts in F and ends in G . Observe that by the design of the algorithm, once the juncture is visited, it splits the message in every angle that intersects T . That is, a juncture is visited in every adjacent face at once.

LEMMA 4.1. *In $MCFR$, for every face F with a visited segment, the segment border node has a message to send across to the unvisited neighbor. A segment-internal node never holds such a message.*

Proof: The proof is by induction on the nodes of a particular face F . A visited segment is created in F when a juncture node is visited. This juncture may be the source or another node splitting the message when it is visited in an adjacent face. Once the visited segment is created, it contains a single border node with two messages sent in the opposite directions. This is our base case.

Let us consider a computation of $MCFR$ where every visited segment of every face is as stated in the conditions of the lemma. First, let us consider a message transmission by node u adjacent to face F . By the induction hypothesis, u may only be a border node. The message recipient v may be an unvisited node or a visited node that has a message for u . Let us consider the unvisited case first. Once unvisited node v receives a message from u , v becomes a new border node with this message while v becomes segment internal without a message. Hence the conditions of the lemma hold. If v is visited, then by definition of the border node, it holds a mate for the message that v transmits. Moreover, by the condition of the lemma, v is also a border node of an adjacent segment. Once v receives the message from u , both mates are destroyed and both u and v become segment-internal nodes. This also preserves the condition of the lemma.

Let us now discuss the transmission of a message by a node u that is not adjacent to F . The only way that it may affect F is if v is a juncture adjacent to F . However, by the design of the algorithm, the juncture is instantly visited in every adjacent face. That is, when v receives a message transmission, it is not visited. Once it receives a message, it creates a new visited segment in F with a single border node v and appropriate outgoing messages.

That is, regardless of the message transmissions, the conditions of the lemma are preserved. \square

LEMMA 4.2. *In $MCFR$, if a face has a visited segment, every node adjacent to this face is eventually visited and none holds messages.*

Proof: If a face with a visited segment contains an unvisited node, then, at least one such unvisited node is adjacent to a border of a visited segment. Due to Lemma 4.1, this border node has a message to be sent to the unvisited adjacent node. Since we only consider fair computations of the routing algorithm, this message is eventually transmitted. Once the message is sent, the adjacent node becomes visited. This process continues until all nodes of the face are visited. Once all nodes are visited they become internal and, according Lemma 4.1, do not hold message. Hence the lemma. \square

LEMMA 4.3. *In $MCFR$, if a face intersects the Steiner tree T , then every node adjacent to this face is eventually visited.*

Proof: Consider the face that contains the source node s . The algorithm starts by creating a single-node visited segment there. According to Lemma 4.2, every node in this face is eventually visited. This includes all junctures adjacent to this face. Repeated application of Lemma 4.2 proves this lemma. \square

PROPOSITION 4.4. *If a target node is connected to the source node, then this node lies on a face that is juncture connected to the source node face.*

The below theorem follows from Proposition 4.4 and Lemma 4.3.

THEOREM 4.5. *Algorithm $MCFR$ guarantees termination and delivery of the message from the source to all targets connected to the source.*




Figure 3: Covering the source-target Steiner hull with unit circles. Illustration for the proof of Theorem 4.7.

Efficiency bounds. Latency of an algorithm is the shortest path that the message may take to reach the target. In multicasting, the latency is the longest such path among all targets. For latency

estimation, following Kuhn et al [19], we assume that the network graph G is bounded degree since such graph can be efficiently obtained in a unit-disk graph by computing a connected-dominating set of G . The latency of a unicast concurrent face routing algorithm is established to be $O(t^2)$ where t is the distance between the source and the target [8, Theorem 2].

THEOREM 4.6. *$MC\mathcal{FR}$ latency is in $O(d^2)$ where d is the Steiner tree diameter.*

Proof: Let d be the diameter of the Steiner tree and m is the number of multicast targets. A Euclidean Steiner tree has at most $m - 2$ virtual nodes [16]. In the worst case, the multicast message in $MC\mathcal{FR}$ has to sequentially reach all nodes in the Steiner tree. That is, the total latency is in $O((2m - 2)d^2)$, which is $O(md^2)$. The theorem's claim follows if the number of targets is constant. \square

Message cost of an algorithm is the total number of messages expended in delivering it to the targets. In the worst case, $MC\mathcal{FR}$ traverses every face of the graph. An edge is adjacent to two faces, hence $MC\mathcal{FR}$ may send up to $2|E|$ messages where E is the number of edges in the graph. However, for most graphs, $MC\mathcal{FR}$ is a lot more efficient. To give a more realistic message cost estimate of $MC\mathcal{FR}$ we make several assumptions about the network graphs.

The graph is *face smooth* if there are two constants c_1 and c_2 that are independent of network parameters such that (i) for each face $p^2 < c_1 a$ where p is the perimeter of the face, and a its area, and (ii) for any two points in the graph, $a_s < c_2 \frac{\pi d^2}{4}$ where a_s is the area of all internal faces that intersect the line between these two points, and d is the Euclidean distance between them. For an internal face, the area computation is straightforward; for the external face, an area of an arbitrary figure enclosing the graph, for example a convex hull, is considered. The first assumption places limits on how "ragged" the perimeter of the face may be, while the second limits how "uneven" the faces may be in size by assuming that the area of all intersecting faces is included in a certain disk whose diameter is related to the distance between two nodes. These assumptions hold for most realistic wireless communication graphs such as unit-disk graphs. *Steiner hull* is the convex hull that contains the nodes of the Steiner tree. It is known that virtual nodes are internal to the Steiner hull [16].

THEOREM 4.7. *For face smooth graphs, the message cost for $MC\mathcal{FR}$ is in $O(|H| + \sqrt{|G|})$, where $|H|$ is the area of the Steiner hull and $|G|$ is the area of the complete graph G .*

Proof: We completely cover the Steiner hull H with unit-disks. See Figure 3 for illustration. In this arrangement, each unit disk covers a square with side length of $\sqrt{2}$. Since k is the maximum node degree, the number of nodes in each unit disk is no more than k . Hence, the number of nodes inside H is:

$$k \frac{|H|}{\sqrt{2}^2} = \frac{k|H|}{2}.$$

Since there are at most k neighbors, each node may be adjacent to at most k edges. A message may be sent across each edge at most twice. Hence, the number messages to be sent inside $|H|$ is $k^2|H|$, which is in $O(|H|)$.

Let us now estimate the number of messages it takes to traverse all the faces that intersect H . Any convex polygon can be inscribed into a rectangle whose area at most twice the size of the polygon [21]. This means that H can be inscribed into a rectangle whose area is at most $2|H|$. The perimeter of this rectangle is $4\sqrt{2|H|}$. We assume that the graph is face-smooth. Combining the two face smoothness conditions we obtain that the sum of the perimeters of all the internal faces that intersect H has this relation:

$$p_s^2 = c_1 c_2 \frac{\pi(4\sqrt{2|H|})^2}{4} = 8\pi c_1 c_2 |H|,$$

which means that p_s is in $O(\sqrt{|H|})$. Similarly, the perimeter of the external face is in $O(\sqrt{|G|})$. Since each face is traversed at most once, the total number of messages used to traverse internal faces that intersect H as well as the external face is in $O(\sqrt{|H|} + \sqrt{|G|})$. Combined with the number of messages needed to traverse faces inside the Steiner hull we get $O(|H| + \sqrt{|H|} + \sqrt{|G|}) = O(|H| + \sqrt{|G|})$. \square

To summarize, Theorem 4.6 shows that the latency of message delivery of $MC\mathcal{FR}$ does not depend on the overall network size, just on the distance between the source and the targets; Theorem 4.7 shows that the total number of messages sent by $MC\mathcal{FR}$ depends on the locations of the targets with respect to the source and the length of the external face of the graph.

5 ABSTRACT SIMULATION

Critical region. Geometric routing is shown to be sensitive to network density [19]. If the network is sparse, then, when the source and destination are connected at all, they are likely to be so close that simple greedy routing is sufficient. If the network is dense, then the presence of a local minimum is unlikely and, again, greedy routing is sufficient. *Critical region* [20] is the range of densities where strategies other than greedy routing are likely to make an impact and determine the efficiency of a geometric routing algorithm. In this range of densities, the route selected by a geometric routing algorithm may significantly deviate from the optimal. Depending on the particular algorithm, the critical region is between 3 and 7 nodes per single node communication area.

Setup. In their classic study of unicast geometric routing algorithms, Kuhn et al [19] use a particular simulation setup to thoroughly evaluate the performance of their algorithms. We extend their setup to use in our simulation.

Specifically, we populate a 10×10 unit square field with nodes placed uniformly at random to achieve a specific network density. The total number n of nodes is equal to the area of the field divided by area of the unit circle and multiplied by the required density d . That is $n = d \frac{100}{\pi}$.

Experiment is a single delivery of a message from a particular source to a particular set of multicast targets. In other words, it is a single complete computation of an algorithm. For each experiment, we generate a new random graph with a randomly selected source and targets. We then calculate each node's neighbors as follows. We first construct a unit-disk graph. For the planar geocasting

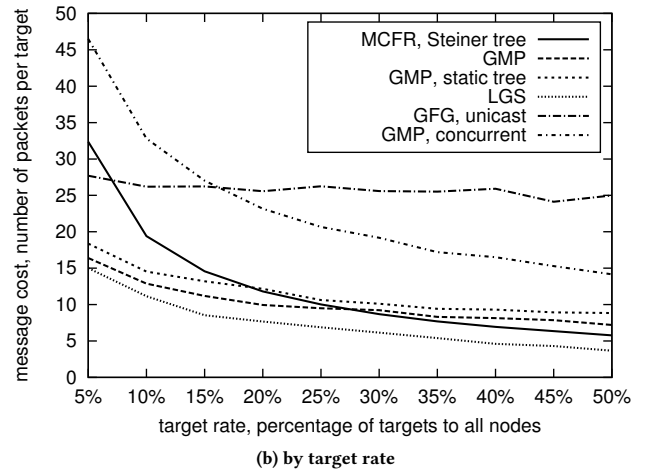

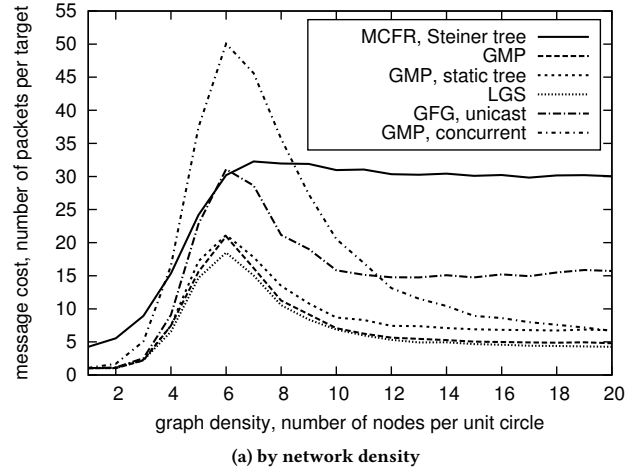



Figure 4: Abstract Simulation: Arrival latency normalized to optimal path (path stretch).

Figure 5: Abstract Simulation: Message cost normalized to number of destinations.

algorithms, we also compute Gabriel subgraph and connected dominating set on it. For each specific data point, we conduct 1000 experiments.

We select the number of destination nodes that achieves a specific *target rate*: the percentage of non-source nodes in the network randomly marked as destinations. For example, target rate of 5% on a density of 7 results in 12 nodes being randomly marked as multicast destinations, rounding up.

We evaluate the message cost and latency of the algorithms. *Message cost* is the number of messages it takes to deliver to all destination nodes. *Latency* is the shortest path taken by the algorithm to reach the destination furthest away from the source normalized to the optimal path to this node.

We estimate both metrics by varying two parameters: network density and target rate. In one experiment, density is varied while the target rate is constant at 5%. In another, density is 7 nodes per unit square while target rate is varied.

Algorithms. We implement both sequential and concurrent multicasting algorithms. For sequential algorithms, we implement: the trivial algorithm that uses a unicast GFG to separately deliver the message to each individual target, LGS [7], GMP as presented in the original paper [30] and a variant of GMP where the Steiner tree is computed once at the source and is not recomputed by every intermediate node. For concurrent multicasting, we implement MCFR, which navigates over Steiner tree. The trees are computed by the source. We also implement a version of GMP where the sequential CFR is replaced with concurrent CFR [8] for message with only a single remaining destination.

Results and analysis. Figures 5 and 4 present the simulation results. Let us discuss the latency results first. The sequential algorithms show a notable peak at the critical region (see Figure 4a). However, the concurrent algorithms show no such peak, and are remarkably close to the optimal path. Target rate has little effect

on latency for any of the algorithms other than the unicast mechanism (see Figure 4b). A slight increase is visible for the sequential algorithms, but none is apparent for the concurrent algorithms. These results show that these algorithms are able to group message delivery without significant increase to any single path within the group.

Let us now discuss message cost evaluation. The majority of the algorithms show a peak at a critical density of about 6 (see Figure 5a). As expected, the GFG unicast algorithm is least efficient of these. MCFR performs the same as the others up until the critical region, after which the cost remains constant rather than dropping as the graph becomes dense. This is due to the lack of any greedy mechanism in MCFR. Target rate does not seemingly affect the GFG unicast algorithm; this is expected. However, it has a notable effect on the other algorithms (see Figure 5b). As the target rate increases, the message cost decreases. We normalize the message cost according to the total number of destinations chosen. Hence, these results show the increase in efficiency of the algorithms as they more effectively bundle messages to groups of destinations.

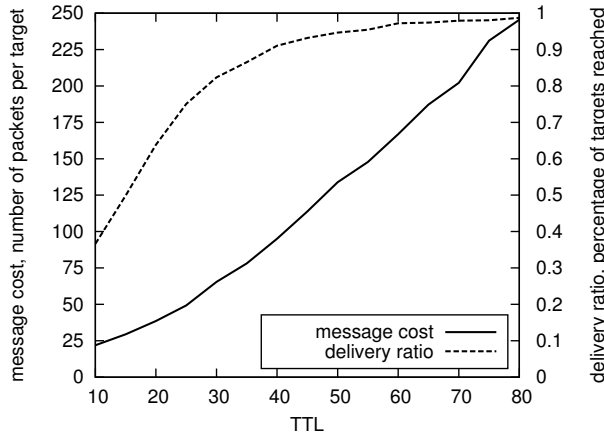



Figure 6: Concrete simulation: Time to live evaluation for MCFR, Steiner Tree, 15 dBm.


6 CONCRETE SIMULATION

Setup. To evaluate the performance of our algorithms, we implement them in WSNNet [4, 10, 11, 13] wireless sensor network simulator. The simulated MAC layer is IEEE 802.15.4 with 866 MHz frequency band and BPSK modulation. The radio model is freespace propagation with constant path loss and rayleigh fading [4]. We simulate 1000×1000 meters field. The unit is 100 meters. The field is populated by nodes placed uniformly at random to achieve a specific network density. To compute the network topology we use a unit-disk graph, then compute a Gabriel subgraph over it. The topology is calculated offline. We evaluate our algorithms' performance at three power levels: 15 dBm, 7 dBm and 0 dBm. The weaker the signal, the less reliable message transmissions are.


As in the abstract simulation, the number of targets is kept at 5% of the total number of nodes and we conduct 1000 experiments for each data point.



(a) 15 dBm



(b) 7 dBm



(c) 0 dBm

Figure 7: Concrete simulation: Delivery ratio.


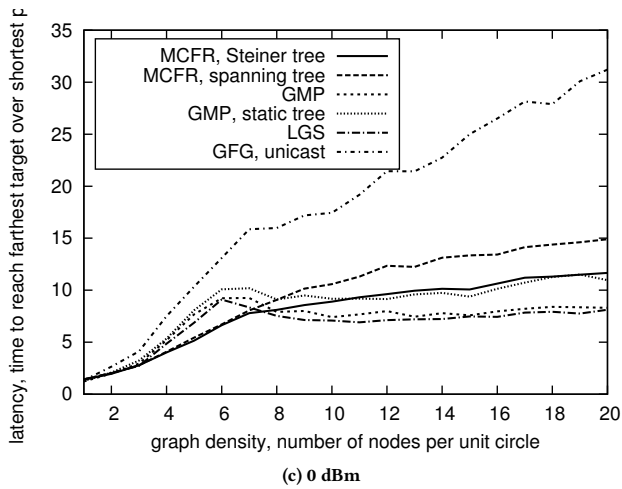


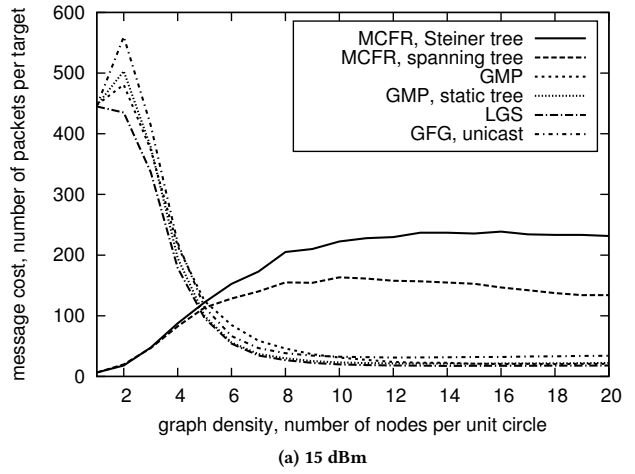



Figure 8: Concrete simulation: Latency.

Figure 9: Concrete simulation: Message cost.

We evaluate the performance of our algorithms according to three metrics. *Delivery ratio* is the number of targets that receive the message divided by the total number of targets. Delivery ratio determines the reliability of the algorithm. *Latency* is the time it takes the algorithm to deliver the message to the target that is geometrically furthest away from the source divided by the time it takes to unicast this message to this target using an optimum route. Latency determines the algorithm's speed of message delivery. *Message cost* is the number of message transmissions divided by the number of targets. Since a message transmission is a radio broadcast, message broadcast to several neighbors is counted as a single transmission. The message cost counts all message transmissions regardless of delivery success. Raw latency and message cost do not take into account delivery success. For example, an algorithm that delivers only to the nearest target may have low latency and message cost. To offset that, we divide the latency and message cost by the delivery ratio.

Algorithms. We simulate LGS, GMP with both source computation and per-node re-computation of the Steiner tree. We simulate two versions of MCFR. In the first version, MCFR navigates over Steiner tree; in the second, it routes over minimum Euclidean spanning tree. The trees are computed by the source.

In realistic environments, both concurrent and sequential algorithms have termination issues. If messages are lost, MCFR packets may not find mates. In this case, the messages may traverse graph faces indefinitely. On the other hand, a sequential algorithm does not detect if a target is disconnected. Again, a message for such disconnected target never reaches its destination. To force termination, we introduce time to live (TTL) for each message. A message is discarded after its TTL expires. To determine optimal TTL, we vary TTL then compute message cost and delivery ratio for MCFR with Steiner Tree and 15 dBm signal strength. The results are shown in Figure 6. The TTL of 55 hops seems to produce the best performance. For the rest of the experiments all algorithms have the TTL of 55.

Results and analysis. The simulation results for delivery ratio, latency and message cost of the simulated algorithms are shown in Figures 7, 8 and 9 respectively. Let us discuss the results. The reliability of MCFR exceeds that of sequential multicasting algorithms. The gap widens as signal strength lowers and message loss increases. For 0 dBm, the delivery ratio of MCFR Steiner Tree never drops below 90% while the delivery ratio of most of the sequential algorithms goes below 40%. Interestingly, the reliability of simple unicasting to all targets exceeds that of GMP and LGS. This is due to the difference in their operation. To optimize message cost, GMP and LGS combine the messages to multiple targets as long as possible. However, a loss of such combined message results in failed delivery to all targets.

The latency of concurrent algorithms is higher than that of sequential algorithms. This is due to the greater number of messages contending for radio channel access. However, with lower signal strength and greater message loss, the gap narrows. For 0 dBm, latency of concurrent and sequential algorithms is similar. The higher delivery ratio of concurrent algorithms match relative speed

of delivery of the sequential algorithms. The only exception is the unicast GFG whose latency continues to be high.

Concurrent and sequential algorithms exhibit different message cost dynamics. At lower densities greater number of targets is disconnected. This forces sequential algorithms to send aimless messages to wander around the network until their TTL expires. Concurrent algorithms do not have this issue. As the network density grows, the relative frugality of the sequential algorithms gives them message cost advantage over concurrent algorithms.

7 CONCLUSION

Algorithm *MCFR* presented in this paper has low theoretical latency and good practical reliability. As future work, we suggest the following. To further improve latency, apply this algorithm to minimum diameter Steiner trees [9]; evaluate fine-grained energy consumption of the algorithms, for example using the approach of Bramas et al [6], as it may impact the survivability of strategic parts of the network.

REFERENCES

- [1] Jordan Adamek, James Scott Robinson, Mikhail Nesterenko, and Sébastien Tixeuil. 2017. Stateless Reliable Geocasting. In *Proceedings of the 36th IEEE International Symposium on Reliable Distributed Systems*. to appear.
- [2] A. Arora, P. Dutta, S. Bapat, V. Kulathumani, H. Zhang, V. Naik, V. Mittal, H. Cao, M. Gouda, Y. Choi, T. Herman, S. Kulkarni, U. Arumugam, M. Nesterenko, A. Vora, and M. Miyashita. 2004. A Line in the sand: a wireless sensor networking for target detection, classification, and tracking. *Computer Networks, Special Issue on Future Advances in Military Communication and Technology* 46, 5 (December 2004), 605–634.
- [3] Sanjeev Arora. 1998. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM (JACM)* 45, 5 (1998), 753–782.
- [4] Elyes Ben Hamida, Guillaume Chelius, and Jean-Marie Gorce. 2008. On the complexity of an accurate and precise performance evaluation of wireless networks using simulations. In *Proceedings of the 11th international symposium on Modeling, analysis and simulation of wireless and mobile systems*. ACM, 395–402.
- [5] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. 2001. Routing with Guaranteed Delivery in Ad Hoc Wireless Networks. *The Journal of Mobile Communication, Computation and Information* 7, 6 (2001), 48–55.
- [6] Quentin Bramas and Sébastien Tixeuil. 2016. Benchmarking Energy-Centric Broadcast Protocols in Wireless Sensor Networks. In *Networked Systems - 4th International Conference, NETYS 2016, Marrakech, Morocco, May 18-20, 2016, Revised Selected Papers (Lecture Notes in Computer Science)*, Parosh Aziz Abdulla and Carole Delporte-Gallet (Eds.), Vol. 9944. Springer, 87–101. DOI: http://dx.doi.org/10.1007/978-3-319-46140-3_7
- [7] Kai Chen and Klara Nahrstedt. 2002. Effective Location-Guided Tree Construction Algorithms for Small Group Multicast in MANET. In *Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Society (INFOCOM-02) (Proceedings IEEE INFOCOM 2002)*, Vol. 3. IEEE Computer Society, Piscataway, NJ, USA, 1180–1189.
- [8] Thomas Clouser, Mark Miyashita, and Mikhail Nesterenko. 2012. Concurrent face traversal for efficient geometric routing. *J. Parallel and Distrib. Comput.* 72, 5 (2012), 627–636.
- [9] Wei Ding and Ke Qiu. 2014. Algorithms for the minimum diameter terminal Steiner tree problem. *Journal of Combinatorial Optimization* 28, 4 (2014), 837–853.
- [10] Tony Ducrocq, Michaël Hauspie, and Nathalie Mitton. 2014. Geographic Routing with Partial Position Information. In *SENSORNETS 2014 - Proceedings of the 3rd International Conference on Sensor Networks, Lisbon, Portugal, 7-9 January, 2014*, Octaviano Postolache, Marten van Sinderen, Falah H. Ali, and César Benavente-Peces (Eds.). SciTePress, 165–172. DOI: <http://dx.doi.org/10.5220/0004872901650172>
- [11] Tony Ducrocq, Michaël Hauspie, Nathalie Mitton, and Sara Pizzi. 2014. On the Impact of Network Topology on Wireless Sensor Networks Performances: Illustration with Geographic Routing. In *28th International Conference on Advanced Information Networking and Applications Workshops, AINA 2014 Workshops, Victoria, BC, Canada, May 13-16, 2014*, Leonard Barolli, Kin Fun Li, Tomoya Enokido, Fatos Xhafa, and Makoto Takizawa (Eds.). IEEE Computer Society, 719–724. DOI: <http://dx.doi.org/10.1109/WAINA.2014.118>
- [12] G.G. Finn. 1987. *Routing and Addressing Problems in Large Metropolitan-Scale Internetworks*. Technical Report ISI/RR-87-180. 87–180 pages.

- [13] Antoine Fraboulet, Guillaume Chelius, and Eric Fleury. 2007. Worldsens: development and prototyping tools for application specific wireless sensors networks. In *Information Processing in Sensor Networks, 2007. IPSN 2007. 6th International Symposium on*. IEEE, 176–185.
- [14] K Ruben Gabriel and Robert R Sokal. 1969. A new statistical approach to geographic variation analysis. *Systematic Biology* 18, 3 (1969), 259–278.
- [15] Michael R Gary and David S Johnson. 1979. Computers and Intractability: A Guide to the Theory of NP-completeness. (1979).
- [16] F.K. Hwang, D.S. Richards, and P. Winter. 1992. *The Steiner Tree Problem*. Annals of Discrete Mathematics, Vol. 53. North-Holland: Elsevier.
- [17] Georgios Karagiannis, Onur Altintas, Eylem Ekici, Geert Heijenk, Boangoat Jarupan, Kenneth Lin, and Timothy Weil. 2011. Vehicular networking: A survey and tutorial on requirements, architectures, challenges, standards and solutions. *IEEE Communications Surveys & Tutorials* 13, 4 (2011), 584–616.
- [18] B. Karp and H.T. Kung. 2000. GPRS: Greedy Perimeter Stateless Routing for Wireless Networks. In *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*. ACM Press, 243–254.
- [19] F. Kuhn, R. Wattenhofer, Y. Zhang, and A. Zollinger. 2003. Geometric ad-hoc Routing: Of Theory and Practice. *22nd ACM Symposium on the Principles of Distributed Computing (PODC)* (July 2003).
- [20] F. Kuhn, R. Wattenhofer, and A. Zollinger. 2003. Worst-Case optimal and average-case efficient geometric ad-hoc routing. In *4th International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*. Annapolis, Maryland, USA, 267–278.
- [21] Marek Lassak. 1993. Approximation of convex bodies by rectangles. *Geometriae Dedicata* 47, 1 (1993), 111–117.
- [22] Jie Lian, Kshirasagar Naik, Yunhao Liu, and Lei Chen. 2006. Virtual surrounding face geocasting with guaranteed message delivery for ad hoc and sensor networks. In *Network Protocols, 2006. ICNP'06. Proceedings of the 2006 14th IEEE International Conference on*. IEEE, 198–207.
- [23] P. C. Liu and R. C. Geldmacher. 1979. On the deletion of nonplanar edges of a graph. In *Proc. of the 10th Southeastern Conference on Combinatorics, Graph Theory, and Computing*. 727–738.
- [24] Martin Mauve, Holger Füßler, Jörg Widmer, and Thomas Lang. 2003. Position-based multicast routing for mobile Ad-hoc networks. *Mobile Computing and Communications Review* 7, 3 (2003), 53–55. <http://doi.acm.org/10.1145/961268.961288>
- [25] Gabriel Robins and Alexander Zelikovsky. 2000. Improved Steiner tree approximation in graphs.. In *SODA*. Citeseer, 770–779.
- [26] Juan A. Sánchez, Pedro M. Ruiz, and Ivan Stojmenovic. 2006. GMR: Geographic Multicast Routing for Wireless Sensor Networks. In *3d IEEE Communication Society Conference on Sensors and Ad Hoc Communications and Networks (SeCon)*. IEEE, 20–29.
- [27] Karim Seada and Ahmed Helmy. 2006. Efficient and robust geocasting protocols for sensor networks. *Computer Communications* 29, 2 (2006), 151–161.
- [28] Godfried T Toussaint. 1980. The relative neighbourhood graph of a finite planar set. *Pattern recognition* 12, 4 (1980), 261–268.
- [29] David M Warme, Pawel Winter, and Martin Zachariasen. 1999. Exact solutions to large-scale plane Steiner tree problems. In *Proceedings of the tenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 979–980.
- [30] Shibo Wu and K. Selcuk Candan. 2006. GMP: Distributed Geographic Multicast Routing in Wireless Sensor Networks. In *26th IEEE International Conference on Distributed Computing Systems (26th ICDCS'06)*. IEEE Computer Society, Lisboa, Portugal, 49.