Strongly-Connected Minimal-Cost Radio-Networks Among Fixed Terminals Using Mobile Relays and Avoiding No-Transmission Zones

Francesco Bernardini[1], Daniel Biediger[1], Ileana Pineda[1], Linda Kleist[2] and Aaron T. Becker[1,2]

***Abstract*— We present strategies for placing a swarm of mobile relays to provide a bi-directional wireless network that connects fixed terminals. No terminals or relays are allowed to transmit into disk-shaped no-transmission zones. We assume a planar environment and that each transmission area is a disk centered at the transmitter. We seek a strongly-connected network between all terminals with minimal total cost, where the cost is the sum area of the transmission disks.**

Results for networks with increasing levels of complexity are provided. The solutions for local networks containing low numbers of relays and terminals are applied to larger networks. For more complex networks, algorithms for a minimum-spanning tree (MST) based procedure are used to reduce the solution cost.

I. Introduction

The problem of providing a connected network, subject to constraints, is closely related to the *minimum range assignment problem for radio networks*. This is a non-deterministic polynomial-time (NP) complete problem [5], where the 2D positions of n terminals are given and the goal is to assign a transmission radius (correlating to a transmission power) to each terminal such that the resulting network is strongly connected, while minimizing the sum of the squared radii. When applied to aerial relays, the problem is called constructing a Flying Ad-hoc Network (FANET) with minimum transmission power [2].

Other related work includes the *relay placement problem* which seeks to place the minimum number of relays to connect a set of stationary terminals. Each terminal is assumed to have a transmission radius of 1 and each relay a radius of r. For this problem, a 3.11-approximation algorithm is given along with a proof that no polynomial-time approximation scheme exists [6].

A necessary condition for a network is that the union of transmission disks must contain all terminals, similar to *minimum-cost coverage of point sets by disks* [1], but this is insufficient to generate a *connected* network.

A network is implied by minimum spanning tree (MST) and Euclidean Steiner tree problems. The Steiner tree is an undirected graph that connects a set of terminal nodes and minimizes the total weight of its edges. It does this by introducing additional internal nodes called *Steiner points*, if this helps reduce the total weight of the edges. In general, these Steiner points have three incident edges, arranged at $120°$ angles. The cost metric for the Steiner tree and the MST is the Euclidean distance of the length of all links in the network. The Steiner tree problem typically generates a network with the minimum length, but for our broadcast model, we use a cost function that increases with the square of the link length. Up to a constant factor, the square link metric is also considered in the *Minimum Area Spanning Tree (MAST)* problem [8], where the area of a tree is given by $\pi/4$ the sum of the squared edge lengths. However, a MAST does not guarantee a strongly connected network and further development is required to obtain the actual range assignment. In Sec. II we mathematically formulate the problem and describe its computational complexity. In Sec. III we provide optimal strategies for two simple scenarios. These solutions inspire the strategies developed in later sections. Section IV defines two heuristic algorithms and applies them to the general problem without and with obstacles, and presents simulation results. Finally, Sec. V summarizes the paper and outlines possible paths forward for this research.

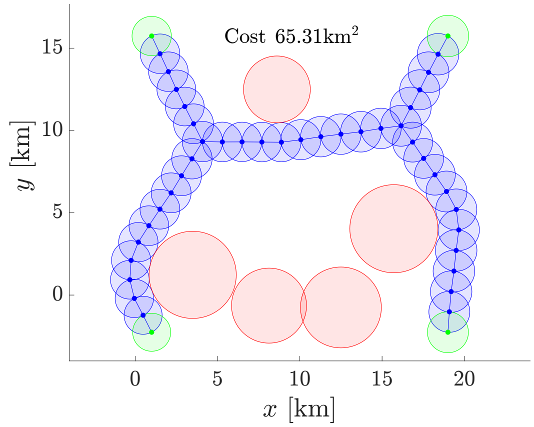

Fig. 1. A network that strongly connects $m = 4$ fixed terminals (in green) by placing $n = 40$ mobile relays (in blue) to establish a strongly-connected mesh network that avoids broadcasting into $\phi = 5$ no-transmission disks (in red). The solution shown minimizes the sum of the terminal and relay broadcast areas (the green and blue disks).

II. Problem Definition

A directed graph is *strongly connected* if every node is reachable from every other node in the graph. In this paper, we assume a node at position A with transmission radius r_A can communicate with a node at position B if the Euclidean distance between A and B is not greater than r_A. A problem

This work was supported by the National Science Foundation under [IIS-1553063, 1932572, 2130793], the Alexander von Humboldt Foundation, and the Army Research Laboratory under Cooperative Agreement Number W911NF-23-2-0014. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

[1] University of Houston, TX USA {fbernardini, debiedig, idpineda, atbecker}@uh.edu

[2] TU Braunschweig, Germany kleist@ibr.cs.tu-bs.de

instance has a set T of m fixed terminals in $\mathbb{R}^{2\times m}$, a set of ϕ obstacles specified by disks with radii R in $\mathbb{R}^{\phi}$ and centers O in $\mathbb{R}^{2\times\phi}$, and a set D of n mobile relays in $\mathbb{R}^{2\times n}$. We then search for a placement of the relays D and an assignment of transmission radii in $\mathbb{R}^{m+n}$ under the constraints that the directed communication graph is strongly connected and none of transmission disks overlap the obstacles. Each transmission radius defines a transmission disk centered at a terminal or relay. See Fig. 1 for a sample solution.

The cost $C(T,O,R,D)$ for a network is

$$C(T,O,R,D) = \sum_{j=1}^{m} r_{T,j}^2 + \sum_{j=1}^{n} r_{D,j}^2, \tag{1}$$

$$\begin{aligned} \text{if} \quad & |O_i - T_j| \geq R_i + r_{T,j} \quad \forall i \in [1,\phi], j \in [0,m] \\ \textbf{and} \quad & |O_i - D_j| \geq R_i + r_{D,j} \quad \forall i \in [1,\phi], j \in [0,n] \\ \text{else} \quad & C(T,O,R,D) = \infty. \end{aligned}$$

Here, $r_{T,j}$ is the transmission radius of terminal j and $r_{D,j}$ is the transmission radius of relay j.

We want to minimize (1). We start by describing solutions to problems with small numbers of terminals, where it is possible to generate optimal solutions. For more complicated instances, we start by computing the MST of the network with squared Euclidean distances as weights to determine the graph topology. To locally optimize the position of relay i, we determine the network neighbors, then move D_i to minimize the required transmission power.

A. Problem Difficulty and Approximation

The range assignment problem of setting transmission powers with fixed transmitters and no obstacles to provide strong connectivity in $\mathbb{R}^2$ is NP-hard [7, Thm 10] and approximating the range assignments in $\mathbb{R}^3$ better than $1+1/50$ is NP-hard [7, Thm 13].

In contrast, there exists a 2-approximation [9] which holds also in 3D. First, compute an MST $\mathcal{T}$ where the weight of each edge is given by the squared Euclidean distances. The cost of $\mathcal{T}$ is a lower bound of the optimal solution for the range assignment problem. Second, each terminal is assigned a transmission range equal to the Euclidean length of the longest incident edge of $\mathcal{T}$. This range assignment ensures strong connectivity of the communication graph between the terminals. As every edge of $\mathcal{T}$ is changed at most once for each incident terminal, this yields a 2-approximation.

Moreover, the analysis is tight as instances of the type shown in Fig. 2 illustrate; the MST solution results in a range assignment where each terminal has radius 1 and, therefore, a total cost of n. The optimal solution has cost $(n/2+1) + (n/2-1)\varepsilon^2$. For $n \to \infty$ and $\varepsilon \to 0$, the ratio goes to 2.

B. Optimal Strategies and Heuristics

By restricting the number of terminals, relays, and obstacles, we can construct problem classes with optimal solutions. We describe two optimal placement cases in the following section. Adding additional terminals and placing obstacles makes the problem harder to compute; Section IV describes heuristic solvers for this problem.

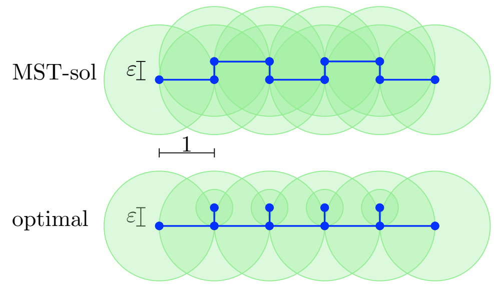

Fig. 2. A 2-approximation using a MST for the range assignment problem. The optimal assignment requires lower transmission power for the offset nodes.

III. Optimal Strategies

A. Two terminals separated by unit radius obstacle

We begin with two terminals, t_1 at $[-d,0]$ and t_2 at $[d,0]$, separated by a unit radius obstacle disk centered at $[0,0]$. Given n mobile relays, what is the lowest cost network according to (1)? Several sample solutions, solved numerically, are shown in Fig. 3.

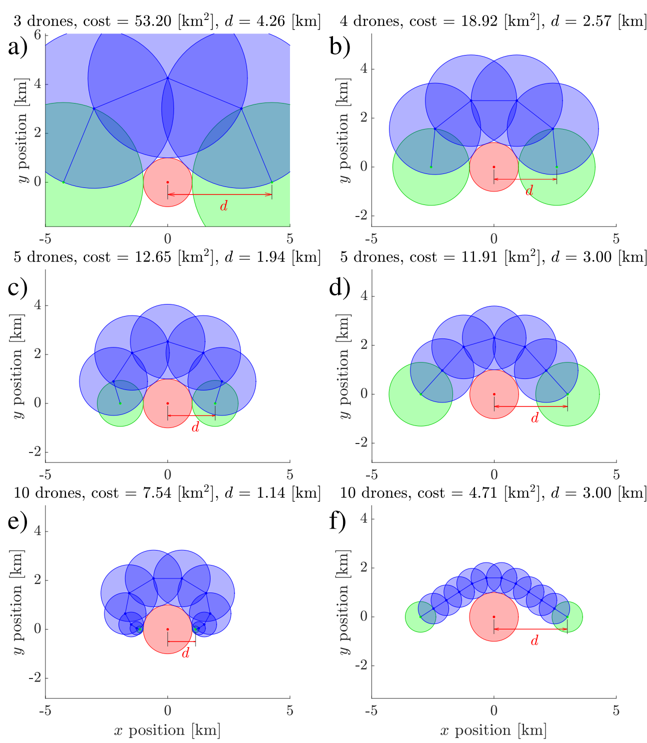

Fig. 3. Building the lowest cost network between two terminals at $[\pm d, 0]$ (in green) separated by a unit radius obstacle (in red), using n relays (in blue). Each of these subplots is shown by a marker in Fig. 4.

The numerical solver adjusts the positions of the n relays and the transmission radii of the relays and the terminals. The transmission radii are constrained to be positive, and the distance from any transmitter to the obstacle must be no less than 1 + the transmission radius. Clearly, an optimal network connecting two terminals forms a single chain.

For a given n, there is a d that minimizes (1). This occurs if all the transmission radii are equal sized and the relays

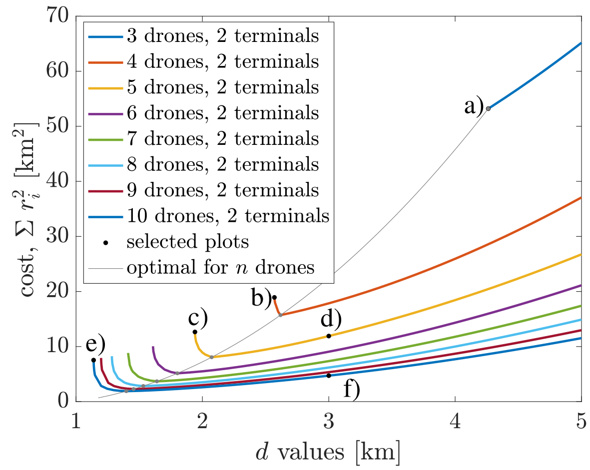

Fig. 4. Building the lowest cost network between two terminals at $[\pm d, 0]$ separated by a unit radius obstacle, using n relays as shown in Fig. 7. Increasing the number of relays n always decreases the cost and decreases the minimum d that can be covered.

evenly distributed on a semicircle of radius

$$d_{\text{min}}(n) = \frac{1}{1 - 2\sin\left(\frac{\pi}{2+2n}\right)}. \tag{2}$$

The relays have angular spacing $\frac{\pi}{n+1}$. As shown in Fig. 4, the plots of cost as a function of d have a minimum at the optimal solution (2) (gray line). For smaller d values the terminal transmission ranges must be less than the optimal value, and the relays' ranges need to be correspondingly larger. For larger d values all the transmission ranges are identical and the path of the relays forms two straight lines that bend in a circular arc about the obstacle. Three or more relays are required for a solution to exist (see Fig. 5).

These optimization problems are relatively easy because we know the communication graph topology *a priori*. When we do not know the communication graph topology before the optimization, the solution must include it as part of the optimization. Even without obstacles, the problem of assigning the minimum area ranges to each terminal is NP-complete [5]. For the simple case of three terminals and one relay, we can find the optimal solution to the problem.

B. Optimal solution: three terminals, one relay

Given a triangle with vertices A, B, C, the relay location D that minimizes the cost for a strongly connected network has multiple candidate solutions, as shown in Figs. 6 and 7.

Theorem 1. *For three terminals at points A, B, C forming the triangle $\triangle(ABC)$, a relay is placed optimally on one of the following three locations:*

1) *the midpoint of the second largest edge of $\triangle(ABC)$,*
2) *1/4 along the perpendicular bisector of the longest edge of $\triangle(ABC)$,*
3) *the circumcenter of $\triangle(ABC)$.*

Proof. Without loss of generality, we assume that $A = [0, 0]$, $B = [1, 0]$, $C = [p, q]$; otherwise we rotate and scale. Let $D = [x, y]$ denote the location of the relay. In an optimal network, the relay is transmitting to two or all three terminals; otherwise it brings no added value.

If D has two neighbors, then it is beneficial to transmit via the shortest side (wlog AC) and place D on the second shortest side (wlog AB). The resulting costs are

$$\begin{aligned}\mathcal{L}_2 = \quad & \underbrace{|AC|^2}_{r_C{}^2} + \underbrace{\max(\{|AD|^2, |AC|^2\})}_{r_A{}^2} \\ & + \underbrace{|BD|^2}_{r_B{}^2} + \underbrace{\max(\{|AD|^2, |BD|^2\})}_{r_D{}^2}.\end{aligned} \tag{3}$$

If $|AD| \geq |AC|$ then $|AD|^2 + |BD|^2$ as well as $\max(\{|AD|^2, |BD|^2\})$ is minimized for $|AD| = |BD|$. Otherwise, we have $|AD| < |AC|$. If $|BD| > |AD|$, then decreasing $|BD|$ and increasing $|AD|$ improves; here we either end in the situation that $|AD| = |BD|$ (claim) or $|AD| = |AC|$ (case 1). Hence, we have $|BD| < |AD| < |AC|$. In this case, $|BD|^2 + \max(\{|AD|^2, |BD|^2\}) = x^2 + (1-x)^2$ is minimized for $x = 1/2$, i.e., $|AD| = |BD|$.

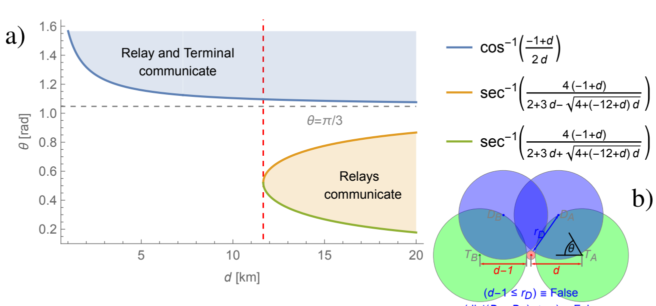

Fig. 5. We require at least three relays to build networks around a unit-radius obstacle with terminals d units on either side of the obstacle as shown in Fig. 3. With two relays placed symmetrically at $[d, 0] + (d-1)[-\cos\theta, \sin\theta]$ and $[-d, 0] + (d-1)[\cos\theta, \sin\theta]$, the transmission disk must not overlap the obstacle disk (schematic b). Schematic a) plots θ values such that the relays communicate in light orange, and θ values such that the relays communicate with their nearest terminal in light blue. These regions have no union, and converge as $d \to \infty$ to $\theta = \pi/3$ (gray dashed line).

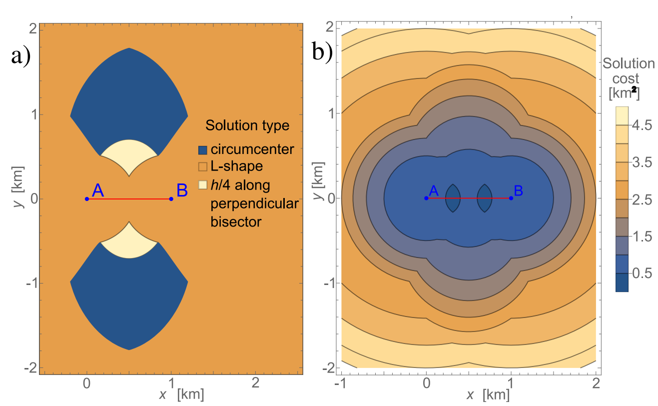

Fig. 6. Given three terminals at A, B, C, there are three candidate solutions for the optimal relay placement to minimize the cost of the network. The solution depends on the shape of the triangle. In the above plots $A = [0, 0]$, $B = [0, 1]$ and $C = [x, y]$.

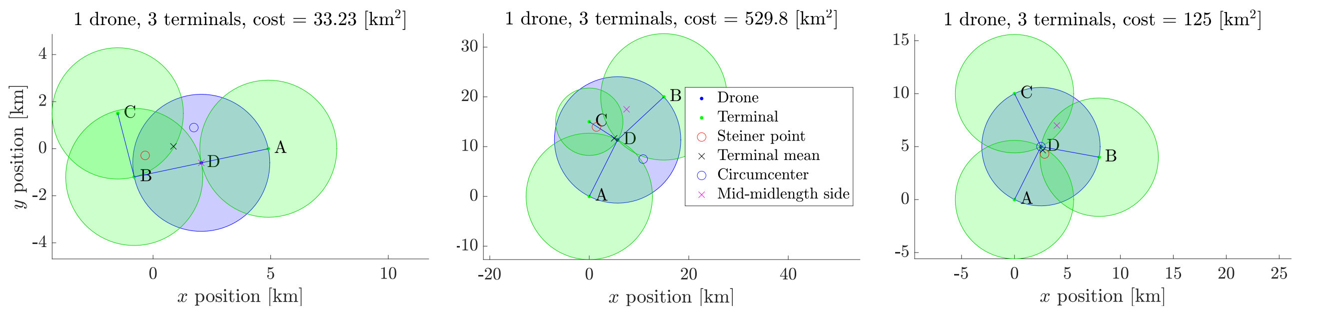

Fig. 7. With three terminals and one relay, there are three types of solutions for the relay position: a) on the midpoint of the second shortest side, b) on the perpendicular bisector of the longest side 1/4 the height of the triangle, or c) at the circumcenter of the triangle.

Now, we consider the case of D having three neighbors. The optimal solution D lies within $\triangle(ABC)$. The resulting costs are

$$\mathcal{L}_3 = \underbrace{|AD|^2}_{r_A{}^2} + \underbrace{|BD|^2}_{r_B{}^2} + \underbrace{|CD|^2}_{r_C{}^2} + \underbrace{\max(\{|AD|^2, |BD|^2, |CD|^2\})}_{r_D{}^2} . \quad (4)$$

We argue that the maximum is attained by two distances; otherwise we may optimize. Suppose the maximum is attained only for $|AD|$; this implies that $x > 1/2$ by $|AD| > |BD|$. We consider the derivative with respect to x, namely $4x+2(x-1)+2(x-p) = 8x-2(1+p)$. If it does not vanish, we may slightly change x and hence improve the cost. If it vanishes, then $x = (1+p)/4 > 1/2$ implying that $p > 1$ and $y \le q/4$ (as D lies in ABC). This yields a contradiction to the fact that $|AD| > |CD|$:

$$\begin{aligned} 16|AD|^2 &= 16(x^2+y^2) \le (1+p)^2 + q^4 \\ &\le (1+p)^2 + q^4 + 8p(p-1) \le (1-3p)^2 + (3q)^2 \\ &< 16(x-p)^2 + 16(y-q)^2 = 16|CD|^2 . \end{aligned}$$

Thus, we conclude that the maximum in eq. (4) is attained by at least two distances. If it is attained by exactly two distances, we may assume without loss of generality, that $|AD| = |BD|$, i.e., D is placed on the perpendicular bisector. Then the cost function simplifies to

$$3|AD| + |CD| = 3(x^2+y^2) + (x-p)^2 + (y-q)^2 ,$$

and its derivative with respect to y reads as $6y + 2(y - q)$ which vanishes exactly if $y = q/4$. If the maximum is attained for all three distances, then $|AD| = |BD| = |CD|$ and D is the circumcenter of the triangle. □

IV. Algorithms for Multiple Terminals and Multiple Obstacles

Solutions to the problem of placing movable relays to enable communications between fixed terminals are explored. We begin by adding obstacles between two terminals and finding a solution strategy. The problem's complexity is increased by adding additional terminals and obstacles.

A. Solving for two terminals, n relays with obstacles

Given two points A and B on a plane, the shortest path that connects them is a straight segment. If the plane contains obstacles, a shortest path that avoids the obstacles may have a different shape.

1) Bitangents method: For simplicity, we assume the obstacles are ϕ circles with centers $\vec{O}_i$ and radii $r_{O,i}$, $i \in \{1, ..., \phi\}$, whose area is excluded from the set of possible coordinates for the points of a path. In this case, the shortest path between two points is given by an alternating sequence of straight-line segments and arcs along the circumference of obstacles [3]. We begin by finding the *bitangents* for each pair of circles, and the tangents from the terminals A and B to each individual circle. If two circles do not overlap, four bitangents exist. These bitangent lines are tangent to both circles. If two circles partially overlap, only bitangents that touch the circles externally exist. However, if one circle is contained inside the other, no bitangents of any kind exist.

Among all the bitangents and tangents determined through this procedure, we keep only those whose line of sight (LoS) is not obstructed by (i.e. do not cross) other obstacles. Next, we add the circular arcs that connect bitangent points on each circle. A pair of points on a circle is always connected by two arcs, which will be different for non antipodal points. If we add to the set of segments the set of all the shortest arcs connecting pairs of points on the circles, we can cast the present system as a graph (V, E, w):

- The terminals, the tangent, and the bitangent points are the nodes V of the graph;
- The tangents, the bitangent segments, and the arcs of circle are the edges E;
- The lengths of each arc or segment are the weights w.

Once the graph is constructed, select a graph search algorithm to determine the shortest path between A and B. Repeating this procedure for every pair of A and B in the accessible domain defines the generalized distance. Figure 8a shows in green the shortest path among A and B built using the bitangent method.

2) k-optimal paths and Yen's algorithm: Introducing obstacles into a simply connected region of the plane generates a connected region where nontrivial *homotopy* classes might exist. Pairs of paths having endpoints in common are said to be homotopic. If there exists a *continuous transformation* that can bring one to the other or vice versa. Intuitively, if the two paths enclose an obstacle, no such continuous function can be defined and the two paths are not homotopic.

In our example of Fig. 8, the green paths in Fig. 8a and Fig. 8b are homotopic, as they do not enclose any obstacle. Conversely, the green paths in Fig. 8b and c are not homotopic to each other or to those in Fig. 8a and b.

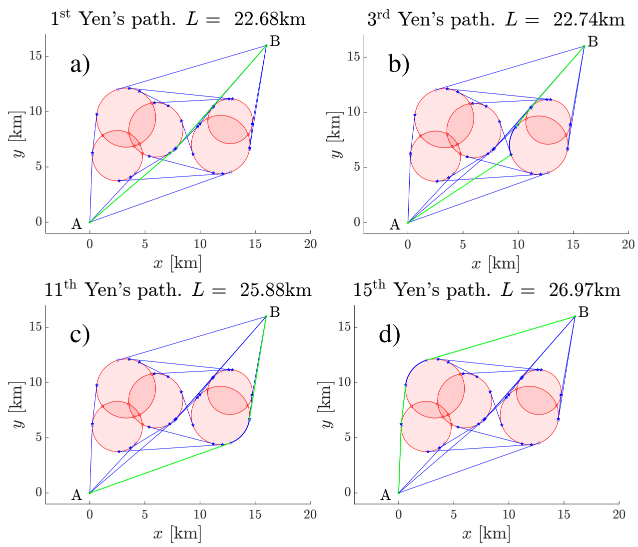

Fig. 8. Construction of Yen's k-optimal paths (green) for a certain configuration of terminals (A and B) and obstacles (red) with the bitangents method. The 1st Yen's path (Fig. 8a) corresponds to the absolute shortest path, particularly, the shortest in its homotopy class. Fig. 8b is in the same homotopy class of Fig. 8a, but it is not the shortest of its class. Fig. 8c and d belong to different homotopy classes and are also not the shortest paths. They are the shortest paths within their respective classes.

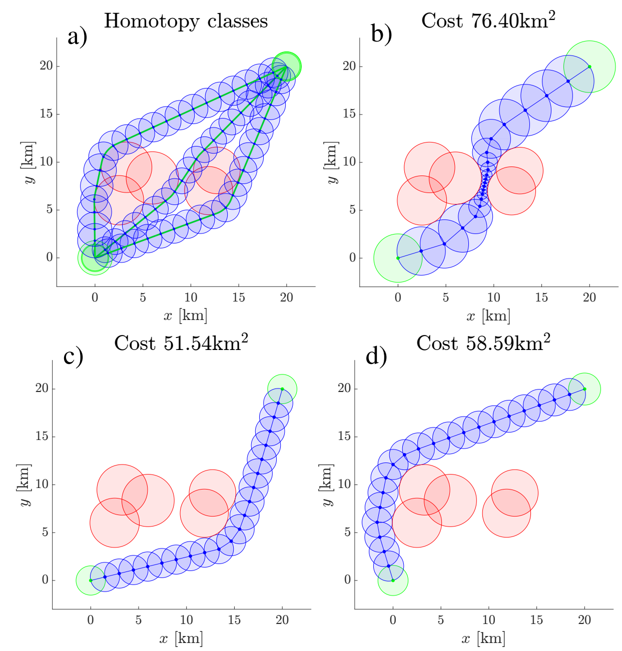

Fig. 9. Homotopy classes (a) for a given configuration of terminals (green) and obstacles (red). Homotopy classes are found by applying Yen's algorithm. A uniform coverage is provided as initial condition using $n = 20$ nodes (blue). A multi-objective optimization (b-d) removes the overlaps between nodes and obstacles seen in a).

Given N obstacles, and assuming paths cannot loop around obstacles, there are at most 2^N classes as each path can have an obstacle to the left or to the right. Paths of different homotopy can be found by considering the $k-$*optimal* paths between A and B, i.e. a set of paths ordered by their length. Using Yen's algorithm [11] applied to the edges of the graph defined in Sec. IV-A.1, such a set can be built. For k sufficiently large all the possible homotopy classes will be visited.

3) 2-dimensional links: The paths found in Sec. IV-A.2 are 1-dimensional. Given two nodes, they are linked by either segments or arcs. Now consider nodes that have circular shape and have a variable transmission radius $r_{D,i}$. Such a radius defines the *coverage* region of the relay. Given n relays, the problem is their optimal placement to minimize the transmission cost between two points A and B .

The system is defined by the set of equations (1), with $m = 2$. This constrains relays from transmitting into obstacle regions. We consider the inter-node distances as a measure of the cost of transmission for each pair and pursue the goal of minimizing this cost. This is a nonlinear multi-objective optimization problem with constraints. We consider a scalarizing function [4] as in (1).

Given an initial number of relays and an initial guess on their positions, the MATLAB function `fmincon` attempts solving the optimization problem while respecting the constraints. Fig. 9 shows the results of the optimization:

- Fig. 9a shows the shortest paths of each homotopy class found in Sec. IV-A.2, where 20 relays have been placed evenly along each class. This is not an acceptable solution of the new problem, as the coverage regions of the relays overlap the obstacles, and so violate the constraints.
- Fig. 9b-d represent the solutions found by `fmincon`, by moving the relays with overlapping discs away from the obstacles, and adjusting the coverage radii to establish the optimal link between each pair of relays.

The existence of a solution depends on the initial number of mobile relays n: if n is too small, the relays will not be able to link the two terminals without overlapping the obstacles. Let n_0 be the initial number of available relays, and L the length of one of the paths in Fig. 9a. Then an average radius R_{avg} can be defined as

$$R_{avg} = \frac{L}{n_0 + 1} \ . \tag{5}$$

If the solution passes through obstacles separated by a distance larger than R_{avg}, (5) is a close estimate of the actual radii of the relays (Fig. 9c-d). Conversely, corridors narrower than R_{avg} result in regions of larger or smaller local densities of relays (Fig. 9b). A new path of length L' may then be covered only if the sum of relays required to cover the different regions is still smaller or equal to n_0.

The minimum cost path in Fig. 9c does not belong to the same homotopy class as the shortest path in Fig. 8. This is due to the choice of (1) as the scalarization function and the additional constraints.

B. Solving for m terminals, n relays, no obstacles

When using more terminals the primary goal remains to connect them through a network at a minimum cost. As in the 1D case, the addition of more relays n affects the nature of the problem: the case $n = 0$ and $n > 0$ generalize the MST and the Steiner tree problem, but with the non-Euclidean cost function (1). This also explains why the solutions need not

Algorithm 1 MST-BASED NETWORK OPTIMIZATION

```
while TRUE do
    Compute MST on set of terminals and relays.
    for relay in relays do
        if relay in a leaf branch of MST then
            Steer relay to parent node with degree > 2.
        else
            Steer relay to average position of neighbors.
        end if
    end for
end while
```

Algorithm 2 ADVANCED NETWORK OPTIMIZATION

```
for node in network do
    Find degree of node.
    if node == terminal and degree == 2 then
        if star improves cost then
            Create star.
        else
            Equilibrate radii.
        end if
    else if degree == 3 then
        Equilibrate radii.
    end if
end for
```

share the properties of the 1D counterparts (e.g. the 120° rule, as shown in Sec. III-B).

With m terminals, the network is a collection of *branches* $\{\mathcal{B}\}$ interconnecting nodes. A branch is defined as a path whose ends are either terminals, or nodes of degree > 2.

The iterative algorithm described in Alg. 1 attempts to solve this problem.

However, this procedure involves local optimization of the cost, and is prone to produce local minima. For example the local rules of the algorithm produce branches with locally uniform densities, but as they do not allow relays to move from a denser branch to a less dense one, densities may not be uniform globally.

To mitigate this issue, assume the system is in a local minimum, and that the radii of the nodes involved in each branch are uniform and equal to $R = L/(N+1)$, where L is the length of the branch and N the number of its nodes *excluding the endpoints*. Let $\mathcal{B}_{\text{sml}}$ and $\mathcal{B}_{\text{lrg}}$ be two neighboring branches (i.e. with one endpoint in common) with higher and lower density of relays and thus **smaller** and **larger** radii.

Then, we move one relay from the branch with the smaller R (R_{sml}) to the branch with the larger R (R_{lrg}), halfway between the branches' common endpoint and the first relay. On the former denser branch the common endpoint is now distant $R' = 2R_{\text{sml}}$ from its nearest neighbor, while on the former less dense side it is distant $R' = R_{\text{lrg}}/2$, both different from their branch average. After the algorithm reaches a new equilibrium the branches will have lengths L' and one more or one less relay:

$$R'_{\text{sml}} = \frac{L'_{\text{sml}}}{(N_{\text{Rsml}}+1)-1}, \quad R'_{\text{lrg}} = \frac{L'_{\text{lrg}}}{(N_{\text{Rlrg}}+1)+1} \,. \tag{6}$$

The cost of each branch is given by $C = (N+1)R^2 = L^2/(N+1)$, therefore, the cost change $\Delta C = C' - C$ is:

$$\Delta C = \frac{L'^2_{\text{sml}}}{N_{\text{Rsml}}} + \frac{L'^2_{\text{lrg}}}{N_{\text{lrg}}+2} - \frac{L^2_{\text{sml}}}{N_{\text{Rsml}}+1} - \frac{L^2_{\text{lrg}}}{N_{\text{Rlrg}}+1} \,, \tag{7}$$

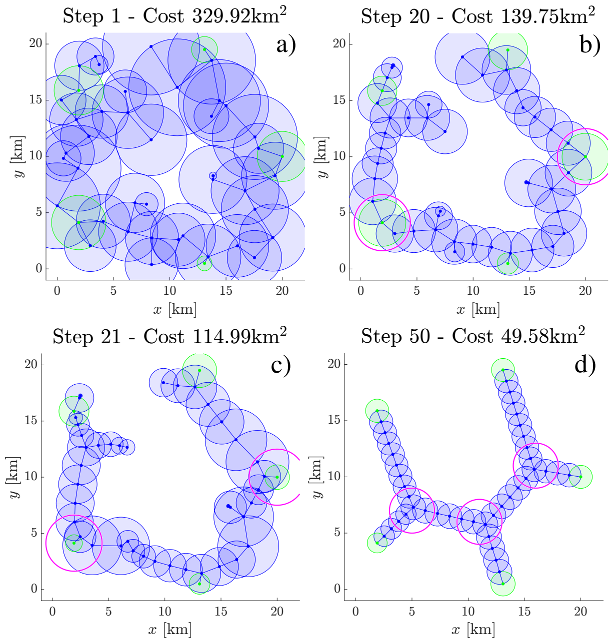

Fig. 10. From an initial random placement of $n = 40$ relays (Fig. a), Algs. 1-2 develop a network interconnecting $m = 5$ terminals. Alg. 2 transforms terminals of degree 2 (Fig. b, circled) into star structures (Fig. c, circled) whenever appropriate, and equilibrates radii across branches. The pseudo-Steiner points (Fig. d, circled) have link angles of $\sim$120°.

Assuming $L_i \sim L'_i$ (reasonable for large N_i) and $L_1 = L_2$, we have

$$\Delta C \equiv \alpha(1 + N_{\text{Rlrg}} - N_{\text{Rsml}}) \leq 0 \,, \tag{8}$$

where α is a positive quantity and the inequality holds for $N_{\text{Rsml}} \geq N_{\text{Rlrg}} + 1$. This happens if the difference between R_{sml} and R_{lrg} is such that $\lceil L_{\text{lrg}}/R_{\text{lrg}} \rceil < \lceil L_{\text{sml}}/R_{\text{sml}} \rceil$. Therefore locally equilibrating radii may result in a cost improvement as long as the density imbalance between two branches is significant and they have comparable length.

Fig. 10 presents the outcome of the optimization under the combined action of Algs. 1–2 at various steps. In particular, Fig. 10d represents the optimal configuration and resembles a full Steiner topology with $m = 5$ terminals and $n = 3$ Steiner points. Due to the high likelihood of local minima, employing only Alg. 1 does not guarantee the result of Fig. 10d. A round of simulations was performed employing only Alg. 1 with $m = 5$ terminals and $N = 100$ different initial placements of $n = 40$ relays: although 66% of the cases had the right topology none of them was optimal. The remaining 44% presented one to three terminals of degree 2 which for this instance are not optimal.

A second round of simulations employed also Alg. 2 and it resulted in 100% of optimal results, up to a rotation of 72°: the terminals are the vertices of a regular pentagon. Figure 10d shows one of the possible solutions while Fig. 11a presents the trends and distributions of costs.

C. Solving for m terminals, n relays, and ϕ obstacles

By combining the concepts introduced in the previous subsections, in particular the heuristics of Algs. 1–2 with

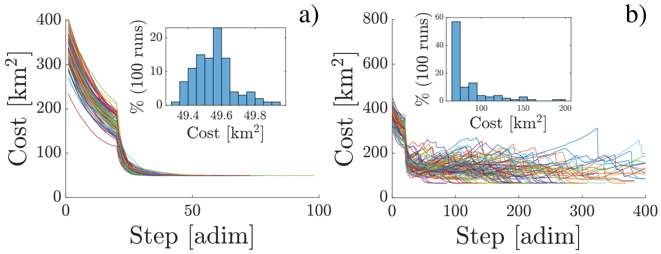

Fig. 11. Cost trend of the 100 simulations of Fig. 10 (a, no obstacles) and Fig. 12 (b, obstacles) using Algs. 1–2. In the case with no obstacles all simulations converged to the optimal configuration of Fig. 10d, up to a rotation of 72° (the terminals are the vertices of a regular pentagon). In the case with obstacles, 53% of the instances obtained the optimal solution of Fig. 12a, whereas the rest either converged to a sub-optimal configuration (21%) similar to those of Fig. 12b-d, or did not converge at all (26%).

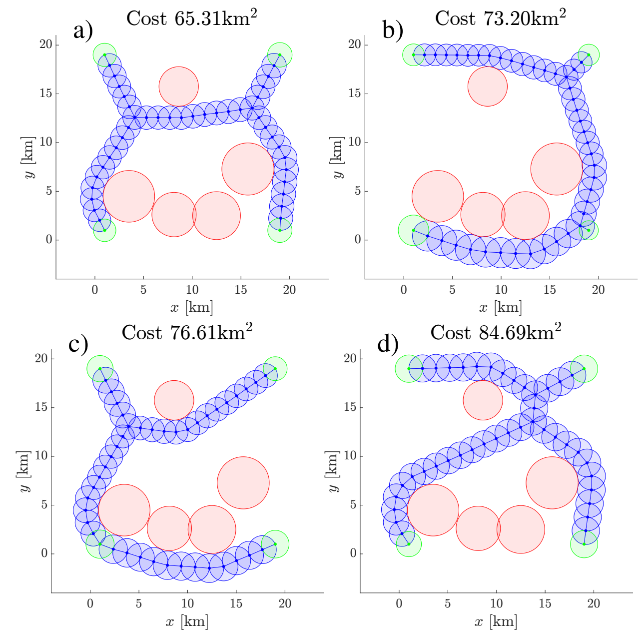

Fig. 12. Converged solutions of four topologies for the same problem set with $m = 4$ terminals, $\phi = 5$ obstacles, and $n = 40$ nodes. The solution that resembles a full Steiner topology has the least cost.

constraint (1), it is possible to ensure that the relays never overlap the obstacle regions.

100 simulations with random initial conditions were performed: the optimal topology was found in only 53% of the cases, while in 21% of the cases the simulation yielded sub-optimal configurations and in 26% of the cases it did not converge at all. Figure 11b presents the trends and distributions of costs for the case with obstacles. Figure 12a shows the optimal configuration while Fig. 12b-d show some of the sub-optimal configurations obtained by the algorithm.

This behavior is imputable to a limitation of the greedy heuristics of Algs. 1–2, which can only overcome local minima related to *local* features, such as density imbalances across adjacent branches. Conversely, obstacles introduce local minima of *topological* nature which, due also to the no-overlap constraint (1), can only be overcome by significant changes in the network structure, incompatible with the greedy heuristics.

Moreover Algs. 1–2 tend to form uniform branches and thus cannot handle narrow passages like those obtained by `fmincon` in Fig. 9b. Fortunately, generating a connection through a narrow passage is expensive, and often cheaper solutions can be found by uniform distributions that route around the obstacles. Nevertheless, these results show that solutions are characterizable by their homotopy classes, i.e. the way they unravel around a given configuration of obstacles. This suggests that the same solutions could be obtained by finding representative networks for each homotopy class and then uniformly placing the relays along them. Thanks to (1) evolution through Algs. 1–2 usually yields a solution in the same homotopy class.

The results of Sec. IV-A.1 can be extended to an arbitrary number of terminals $m > 2$. First, the graph built using terminals and tangents/bitangents is extended to include edges among each pair of nodes in LoS. Then, this graph is used to set up and solve a Steiner tree problem for graphs (STPG), obtaining a solution with shortest length. This solution will be the shortest representative of its homotopy class. To obtain representatives of other homotopy classes the graph is altered by removing edges and recomputing the solution.

An application of these concepts is shown in Fig. 13a. To solve the STPG we implemented SCIP-Jack [10] into custom MATLAB routines. The solution of the STPG (Fig. 13b) is then used as an educated guess for the initial uniform placement of the relays (Fig. 13c). Although optimal according to the Euclidean norm, this choice requires the relays to pass through corridors much smaller than the average radius of the network (Fig. 13d). This will likely result in higher costs if a solver like `fmincon` is used, or a failure if the rules of Algs. 1–2 are applied, so it is discarded and the edges passing through the obstacles are removed. A new solution is determined (Fig. 13d) and the same procedure is used to position the relays (Fig. 13e). Finally Algs. 1–2 are applied and the a feasible configuration is obtained (Fig. 13f).

This result generalizes the conclusions of Sec. IV-A.3. The shortest tree and the optimal range assignment solutions in the presence of obstacles can belong to different homotopy classes. This approach is a heuristic extension of Yen's algorithm, valid for paths, to trees.

D. Estimating convergence likelihood

Given n relays, the convergence likelihood for an initial configuration under Algs. 1–2 can be estimated as follows: let L be the length of the solution of the STPG, then a preliminary placement of n relays along its branches would give an average radius $R_{\text{max}} \simeq L/n$, analogous to (5). If the solution passes between pairs of obstacles whose inter-obstacle distance

$$D_{OO'} \equiv \|O - O'\| - R_O - R_{O'} \ , \tag{9}$$

is smaller than R_{max} then the solution is likely to either be suboptimal or not converge under Algs. 1–2. A similar argument holds also for terminal-obstacle distances:

$$D_{TO} \equiv \|T - O\| - R_O \ . \tag{10}$$

The solution of Fig. 13b has $L = 50.59\,\text{km}$ and $n = 40$ so $R_{\text{avg}} \simeq 1.26\,\text{km}$ while the smallest inter-obstacle distances are $D_{34} = 0.07\,\text{km}$ and $D_{25} = 0.63\,\text{km}$, both smaller

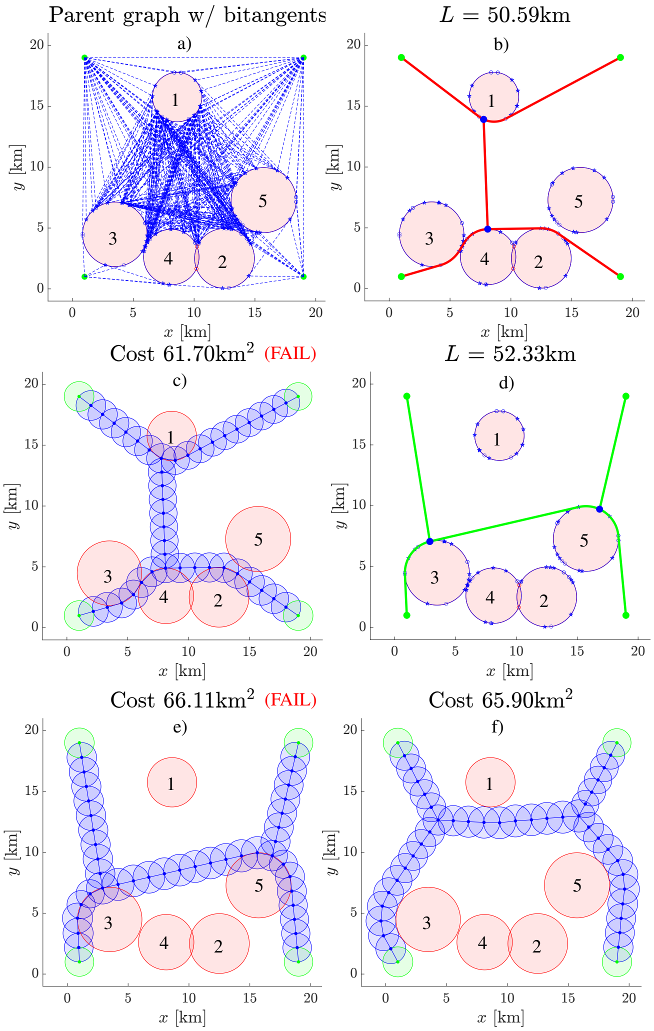

Fig. 13. An STPG can be paired with Algs. 1–2 to determine a better initial condition for the relays. The parent graph is obtained by connecting all vertices in LoS (a). The solution of the STPG is obtained (b). Relays are placed uniformly along the solution, but evolution may not converge if narrow passages between obstacles are involved (c). Sub-optimal solutions which avoid narrow passages are found in a different homotopy class (d). Relays are placed uniformly along the solution (e). The system is evolved with Algs. 1–2 until equilibrium (f). Configurations c) and e) are intermediate states that violate the no-overlap constraint (1).

than R_{avg}. As expected, the simulation did not converge. Removing graph edges that enable passage between obstacles 3–4 and 2–5, the solver may now determine sub-optimal networks such as Fig. 13d. This network passes through obstacles 1–3 and 1–5 which have $D_{13} = 7.69\,\text{km}$ and $D_{15} = 6.34\,\text{km}$, both larger than R_{max} and, as expected, the simulation converged.

V. Conclusions and Paths Forward

This paper describes a framework to solve the minimum range assignment problem amidst obstacles, assigning to nodes of the network radii equal to the largest incident edge, and assuming a cost function quadratic in these radii. The network is optimized with the constraint that network nodes and obstacles, both modeled as disks, do not overlap. The heuristics defined in Algs. 1–2 were first tested without obstacles, where they proved to be independent from initial conditions and yielded the optimal result in 100% of the trials for a specific configuration. Introducing obstacles spoiled this independence: trials converged to several stable configurations, but only found the optimal configuration in 53% of the trials. These configurations are topological local minima which cannot be overcome by the heuristic rules of Algs. 1–2. A procedure was devised to systematically explore homotopically different networks and use them as guidance for initial relay placement. Evolution under Algs. 1–2 can then be biased to produce solutions of the range assignment problem in the same homotopy class of the initial placement. The optimal Euclidean length solution and the optimal range assignment solution are not necessarily in the same homotopy class. This underscores the relevance of homotopy class in the $m > 2$ case and suggests the usefulness of a formal extension of Yen's algorithm from paths to trees.

Extending these results to three dimensions with spherical obstacles, and to non-convex terminal geometries are avenues for future research. Natural applications of these ideas are in distributed networks communicating via line of sight (1D) or coverage (2D) methods: for example, a network might be composed of unmanned aerial vehicles (UAVs), ground/underwater remotely operated vehicles (ROVs), sensors, or satellites.

References

[1] H. Alt, E. M. Arkin, H. Brönnimann, J. Erickson, S. P. Fekete, C. Knauer, J. Lenchner, J. S. Mitchell, and K. Whittlesey, "Minimum-cost coverage of point sets by disks," in *Proceedings of the twenty-second annual symposium on Computational geometry*, 2006, pp. 449–458.

[2] A. Altaweel, H. Mukkath, and I. Kamel, "GPS spoofing attacks in FANETs: A systematic literature review," *IEEE Access*, pp. 1–1, 2023.

[3] E.-C. Chang, S. W. Choi, D. Kwon, H. Park, and C. K. Yap, "Shortest path amidst disc obstacles is computable," in *Proceedings of the twenty-first annual symposium on Computational geometry*, 2005, pp. 116–125.

[4] T. Chugh, "Scalarizing functions in bayesian multiobjective optimization," in *2020 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, 2020, pp. 1–8.

[5] A. E. Clementi, P. Penna, and R. Silvestri, "On the power assignment problem in radio networks," *Mobile Networks and Applications*, vol. 9, pp. 125–140, 2004.

[6] A. Efrat, S. P. Fekete, P. R. Gaddehosur, J. S. Mitchell, V. Polishchuk, and J. Suomela, "Improved approximation algorithms for relay placement," in *Algorithms-ESA 2008: 16th Annual European Symposium, Karlsruhe, Germany, September 15-17, 2008. Proceedings 16*. Springer, 2008, pp. 356–367.

[7] B. Fuchs, "On the hardness of range assignment problems," *Networks: An International Journal*, vol. 52, no. 4, pp. 183–195, 2008.

[8] D. A. Guimarães and A. Salles da Cunha, "The minimum area spanning tree problem: Formulations, benders decomposition and branch-and-cut algorithms," *Computational Geometry*, vol. 97, p. 101771, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0925772121000274

[9] L. M. Kirousis, E. Kranakis, D. Krizanc, and A. Pelc, "Power consumption in packet radio networks," *Theoretical Computer Science*, vol. 243, no. 1-2, pp. 289–305, 2000.

[10] D. Rehfeldt and T. Koch, "Implications, conflicts, and reductions for Steiner trees," *Mathematical Programming*, vol. 197, pp. 903 – 966, 2023.

[11] J. Y. Yen, "An algorithm for finding shortest routes from all source nodes to a given destination in general networks," *Quarterly of applied mathematics*, vol. 27, no. 4, pp. 526–530, 1970.