Optimal Placement of Radio Repeater Networks

Control #10754

February 14, 2011

Abstract

In this paper we consider the problem of placing radio repeaters to serve users in an area. Given a population distribution and geographical map, we use a hill-climbing algorithm to find a minimum number of repeaters required to cover an area and then a genetic algorithm to provide maximal population coverage and network connectedness. We then use hill-climbing techniques to allocate subnetworks based on population size at repeater locations so that two arbitrary users can communicate even when all other users are communicating over the maximum number of possible networks. Our resulting algorithm is capable of producing a range of repeater network allocations, from robust networks that are capable of handling worst-case usage scenarios to smaller networks that provide optimal population coverage and connectivity. On a set of real-world population and geography data, we found that the combination of the hill-climbing and genetic algorithms had 28% better population coverage than a control algorithm did, as well as higher connectivity.

Contents

List of Figures

1 Overview

Radio communications today are ubiquitous, which makes the associated problems that arise in enabling radio communications more important. Eighty years ago, radios were bulky devices capable of sending out and listening to waves on set frequencies. Today most wireless communication is conducted through infrastructure that allows users to communicate with other specific users around the world via satellites and towers in which the older problems of radio communication are handled digitally. Local communication, of the sort utilized, say, by a police or fire department, still relies on the cheap and reliable VHF (Very High Frequency) or UHF (Ultra High Frequency) spectrum.

In this paper we consider a scenario in which a very large number of people are in need of communicating over the VHF spectrum at once. Perhaps a natural disaster has occurred, rendering cell phone and Internet connections unusable. Perhaps there is a war zone in which it is difficult to maintain communications via more modern methods. The solution to such scenarios is to put into place the infrastructure for local communications using VHF.

In section 2 we outline the important features that a model should consider . Our model has been split into two parts, each solving a different problem as discussed in section 3. The complete model is presented in section 4.

2 General observations

In this section we discuss observations of the problem statement, ambiguities, and general assumptions we have made to simplify and clarify the modeling process. Additional clarification of some assumptions is provided in the appendices.

2.1 Initial considerations

Signal strength: Each user of the system has a radio which they use to communicate. The radio can be tuned to different frequencies, however, the signal strength of the radio is too weak to communicate over long distances. This problem is addressed using repeaters, which are devices that listen in on a certain frequency and then re-broadcast the signal they've received at a different frequency and with much greater strength. Signal strength is not a simple function of distance, but varies with elevation as well.

Interference: The VHF spectrum of frequencies only allows for a certain number of simultaneous users before they begin to interfere with each other. This issue is addressed by implementing sub-audible tones which users broadcast along with their signal. Certain repeaters are then tuned to only re-broadcast signals which contain a certain sub-audible tone. This allows multiple repeaters to be placed in the same location with the same broadcast frequency.

Networks: By tuning repeaters to certain frequencies and sub-audible tones and placing them within range of each other, it is possible to create repeater networks which cover larger areas.

Usage: In any given area there are some users who wish to communicate with some other area. Since a repeater or repeater network can only service one conversation at once (although one conversation can have multiple participants) this necessitates the placement of multiple repeaters or repeater networks in certain areas. In addition, certain networks must be made available in designated areas, such as within police jurisdictions. A naive approach would be to place repeaters number of repeaters in an area based solely on population density, however, this is only one aspect of the problem, as high population density areas separated by low population density areas will definitely want to communicate between each other, necessitating the addition of repeater networks in low population density areas.

Cost: Repeaters vary in cost depending on range, features, location, and elevation. So it is important to place as few repeaters as possible in order to achieve a given objective.

2.2 Simplifying assumptions

We made the following simplifying assumptions for creating this model.

Network type: The technical ways in which repeaters interact are complex. It is not an easy task to construct a viable network of repeaters that works correctly as each repeater must be tuned to the right frequency and sub-audible tone to communicate with other repeaters. In reality, however, we observe a large variety of so-called 'half-duplex' network types, or networks in which it is possible for two-way conversations where only one user can talk at a time. So for the purposes of this model we assume that we can build arbitrarily complex networks and we ignore the technical details of how they link to each other. The exception is that we impose an upper bound for the possible number of functioning networks in one area since eventually interference becomes an issue. For a more thorough justification of this assumption, please see appendix A.

Repeater connections: Repeaters must be within a certain range to communicate with each other. The range with which they can communicate with each other is greater than the range with which users can communicate with the repeater. So the problem of building connected repeater networks in an area is solved by solving the problem of covering an area with enough repeaters so that a user can broadcast to and listen to a repeater when they are in the area as long as the user's maximum radius of communication is less then one half of the repeater's maximum radius of communication. For a further discussion of this point please see appendix B.

Repeater redundancy: One issue which we might have considered is the redundancy of repeater networks. Repeaters might fail, in which case it might be nice to have nearby repeaters to re-route communications We did not consider this in our model, instead, we assumed that it would be cheapest and therefore best to mount multiple repeaters in the same location, if possible.

Repeater height: The cost and the range of the repeater depends on the height which repeaters are placed at. However, for our model we pick a standard height, usually about 100-200 feet above the ground, at which the repeaters must be mounted.

Area under consideration: For the more complicated model we assume that the area we are covering is rectangular. Although the problem statement asks for a circular area, our model allows for circular areas as a special case in which the circle is inscribed within a rectangle and the area outside of that circle is empty.

3 Elements of the model: Simple case, geography, then population

In this section we discuss our model, beginning with a simple case and adding additional elements.

3.1 A very simple case

We begin with a circular area C of some radius R. It is completely flat and has a population of 2 users who wish to communicate with each other. Each user can be heard by and pick up signals of repeaters at a distance of r. The problem is to place the minimum number of repeaters such that the entire map is covered by circles of radius r centered at each repeater.

Clearly if $r = R$ then the optimal solution is to place 1 repeater in the center of the map.

If $r < R$ then 2 repeaters will not be enough no matter where they are placed, because each repeater will only be able to cover less than half of C's circumference and there will always be points on the circle which are not covered. However, three repeaters will be enough for r close to R.

However, once $r < 3R/4$ three repeaters will not be enough. 4 will not be enough either. Perhaps 5 will suffice, as in figure 1.

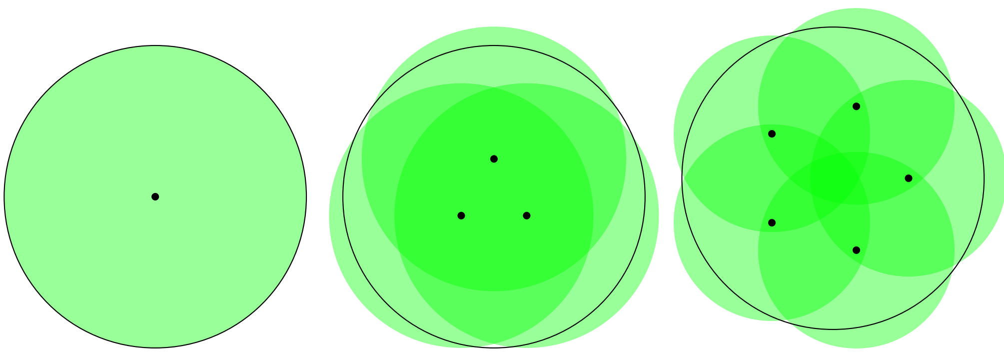

(a) $r \geq R$ (b) $r < R$ (c) $r < 3R/4$

Figure 1: Various user radii

In general, the limit of the ratio of the total area to the area each repeater area as the radius goes to zero is

$$\lim_{r \to 0^+} \frac{1}{r^2 N(r)} = \frac{3\sqrt{3}}{2\pi},$$

where $N(r)$ is the total number of discs required to cover the entire circle [?].

Now imagine that there are 1,000 users instead of two. Each user is paired with another and they all wish to speak with each other. Then for some given radius of the user's range it's clear that the best approach is to take the most efficient covering of the area with repeaters and duplicate it 499 times in order to accommodate all 500 users. Unfortunately, there are a limited number of frequencies and subtones available, so this may not be possible.

This simple thought experiment, while not directly applicable to our main model, illustrates some of the important features of repeater placement.

- The problem is not always solvable. It may be that there are too many users who want to communicate to too many other users, and since there are a limited number of frequencies and sub-audible tones which repeaters can operate at without interference it may not be possible to service all the users.
- As the size of the area increases the number of repeaters required to form a network which completely covers the area increases rapidly.

3.2 Introducing geography

Now we introduce a wrinkle - geography. By geography we mean both the positions of things like rivers and lakes, and also the elevation of a point above sea level.

3.2.1 Problem statement

Let $D = [0, L_x] \times [0, L_y]$ be a set of ordered pairs. Let $T : [0, L_x] \times [0, L_y] \to \{0, 1, 2\}$ be a function. Let $d : [0, L_x] \times [0, L_y] \to \mathbb{R}_+$. Then the problem is to find a set $R = \{r_1, r_2, \ldots, r_M\}$, called the *allocation*, such that given any point $a \in D$ with $T(a) \neq 0$ there exists some r_i such that $d(a, r_i) < K$, where K is some constant. The set R must also satisfy the property that $T(r_i) \neq 0, 1$. Furthermore, if we consider the points of R to form a graph (R, e) where e is the set of all r_i, r_j such that $0 < d(r_i, r_j) < K$ for all $i \neq j$ then (R, e) must be a connected graph.

In this problem D represents the area under consideration, The function T represents the type of land at a point: 0 is empty space, 1 is water, and 2 is land. The function d is a distance metric between points on the map. The elements of the set R represent radio repeater locations, and the problem is to find a distribution of radio repeaters that gives coverage to every point on the map while minimizing the number of repeaters. Note, however, that we assumed areas covered in water were equally important for radio coverage as land.

3.2.2 Choosing d to incorporate elevation

This model has the convenient feature that implementing elevation is an easy task. For initial testing we defined d to be the Euclidean distance between two points, i.e., given points (x, y) and (x', y') we defined

$$d = \sqrt{(x - x')^2 + (y - y')^2}.$$

In order to model the effects of elevation we used the Longley-Rice Irregular Terrain Model[4] (LRIT) as implemented in the open-source program SPLAT![5] to measure how much a signal would degrade between two points. Then we defined

$$d(a, b) = \frac{2}{1 + \deg(a, b)} - 1,$$

i.e., the reciprocal of how much a signal transmitted at a would have degraded by the time it reached b as computed by SPLAT!. It is possible in LRIT for a signal to degrade completely, in which case the distance between the two points is 1. If $\deg(a, b) = 1$, which means the signal has not at all degraded, then the distance between the two points is 0. For the model we would then choose as a value of K something like $\frac{1}{3}$ to ensure that signals between repeaters have not degraded any more than $\frac{1}{2}$ their original strength.

3.2.3 Solution

Solving the model posed significant challenges, as depending on our choice of d and the given function T the problem is not even necessarily solvable. For example, consider the terrain in figure 2.

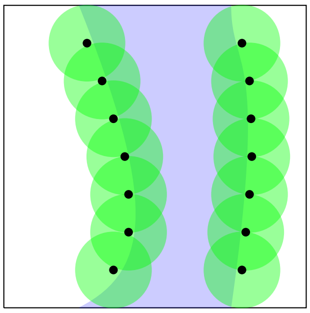

Figure 2: There is no way to connect either shore to the other across this vast river

We implemented two optimization techniques [7] [6] in order to find near-optimal solutions. First, a hill climbing algorithm was used to find an upper bound on number of repeaters needed to completely cover the area. We generated a large random allocation of repeaters across the map such that all points on the map were covered. A greedy heuristic was then used for each repeater. If it could move to a feasible spot where it covered more area while covering the same area as it did before then it did. Any repeaters which could be removed without uncovering some spot were then removed and the process was repeated.

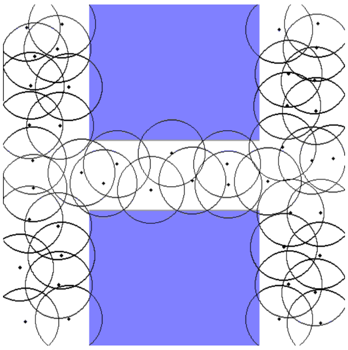

(a) Hill climbing

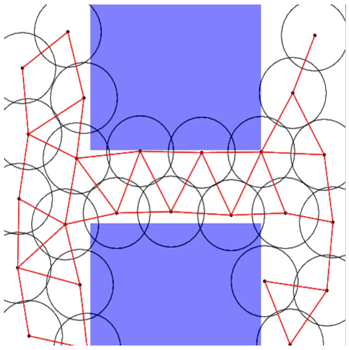

(b) Genetic

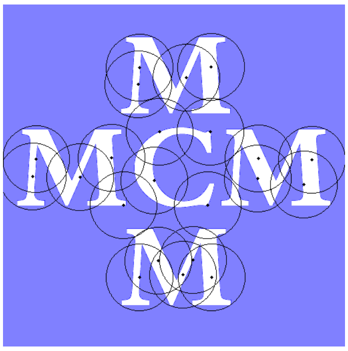

(c) Hill climbing

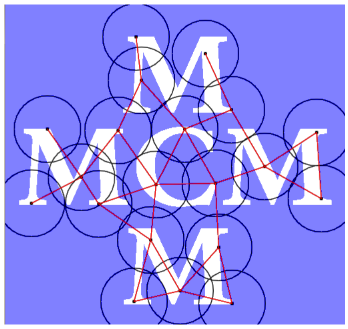

(d) Genetic

Figure 3: Both algorithms

This algorithm tended to produce too many repeaters, as shown in figure 3a. We thus implemented a genetic programming technique as in [2] to refine the results. The genetic algorithm was implemented as follows: A number (100) of random feasible allocations of repeaters of some fixed size M was generated. We chose M to be the number of repeaters which the hill climbing algorithm found could cover the entire area. A fitness heuristic was computed which took the following factors into account:

- Area covered: The more area within range of a repeater the higher the fitness. However, areas which are double-covered are penalized slightly.
- Connectedness of the graph: The more disjoint subgraphs the lower the fitness.
- Population density: Double-covering areas is penalized, but less so for double-covering areas with population.

The algorithm then repeats but with size $M - 1$ of repeaters to allocate until areas of land start to lack coverage. This algorithm tended to produce nice results, as seen in figure .

Due to the heavy computational requirements of implementing elevation we did not actually compute a repeater distribution which took elevation into account. However, after measuring the computational requirements of LRIT and our algorithm we estimate that it would take about two days to compute on an average desktop computer. Since repeater distributions are not required to be computed in real-time this is an acceptable feature of the model.

3.3 Introducing population

Now we consider population distribution in the model. The chief problem here is that it only takes one pair of people to completely tie up a half-duplex repeater network. With our previous problem we tried to efficiently cover the map. If we wanted to serve 1000 people with what we have so far the straightforward approach would be to add 499 duplicate repeaters tuned to different sub-audible tones at each location. Then in the worst case scenario, that there are 500 pairs of individuals looking to talk to each other, the network would be able to serve every user. Unfortunately, this would be prohibitively expensive and could not work as there would be far less than 499 possible frequency/sub-audible tones available. In order to find a cheaper and more workable solution we introduce population data into the model. The objective for this model will be to create smaller repeater networks which serve specific areas.

3.3.1 Problem statement

Let $R = (V, e)$ be a connected graph where $V = \{r_1, r_2, \ldots, r_M\}$. Let $p : V \to \mathbb{N}$ be a function, and let $\psi : R \to [0, 1]$ be some other function. Then we want to find a collection of subgraphs of R, called the *allocation* of repeater networks, $\mathcal{R} = \{R_1, R_2, \ldots, R_N\}$, such that

1. Each subgraph is also connected.

2. Given any two r_i, r_j, if $\mathcal{R}'$ is the set of elements of $\mathcal{R}$ which connect r_i to r_j, R' is the set of elements of R which are vertices of any element in $\mathcal{R}'$, and $\varphi(r_i, \mathcal{R}')$ is the number of elements of $\mathcal{R}'$ which r_i is a vertex of, then

$$|\mathcal{R}'| \geq \min\{p(r_i), p(r_j)\} + \left(\frac{1}{2} \sum_{r_k \in R', r_k \neq r_i, r_j} \min\{p(r_k), \varphi(r_k, \mathcal{R}')\}\right) \psi(r_i, r_j).$$

3. If $\mathcal{C} = \{C_1, \ldots, C_{N'}\}$ is some other set of the subsets of R with these two properties then
$$|C_1| + \ldots + |C_{N'}| \geq |R_1| + \ldots + |R_N|.$$

4. The vertex r_i shows up at most S times in the set of subgraphs, where S is some fixed number.

This problem is analogous to finding a number of subnetworks to best serve the population of the map, which here means to find an allocation of subnetworks such that any two users can be guarenteed to be able to contact each other even in the case where all the other users are communicating in the `worst' possible way. The graph R represents the possible repeater locations. The function p represents the total population served by a repeater location. The function ψ is some discounting function, possibly based on distance.

Condition 1 requires that each subgraph correspond to a repeater subnetwork. Condition 3 enforces optimality of the solution.

Condition 4 restricts the possible number of repeaters at a point by some number, presumably the number of possible sub-audible tones available.

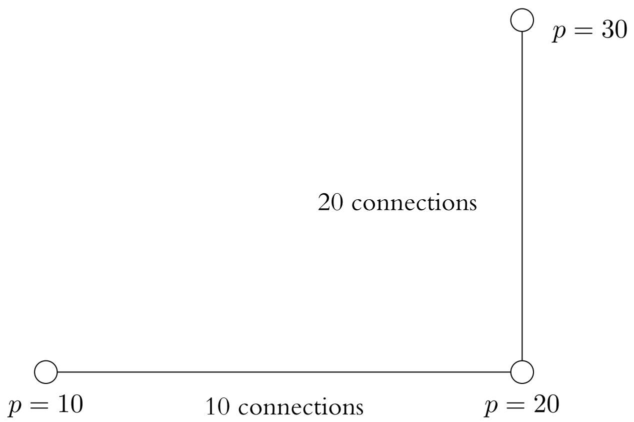

Figure 4: The maximum number of connections between two points is the minimum population between the two points

Condition 2 restricts the possible repeater networks by requiring that two points be able to communicate even if other networks are tied up. The minimum of the population between two points was chosen because that represents the maximum number of connections which might be made between two repeater locations, as seen in figure 4. However, we also need to adjust for the fact that each repeater might be used by a pair of people, which explains the second term in condition 4. If a repeater network covers 6 repeater locations, then in order to serve every person in the area in the worst case scenario of each person communicating with one other person we will need at least 3 repeater networks. However, if the population at a location r_k is less than the number of subnetworks which pass through that point, $\varphi(r_k, \mathcal{R}')$, then we only count people until the population runs out, which explains why there is a second minimum in the expression.

Finally, the function ψ is added in order to possibly reduce the number of repeaters if we want to discount for distance. Note that the addition of the function ψ guarentees that a solution exists for some form of ψ, for example $\psi \equiv 0$.

3.3.2 Choosing ψ to reduce number of repeaters

If $\psi = 1$ then the problem becomes finding a distribution of networks such that an arbitrary pair of users will be able to communicate even if every other user is communicating using the maximum number of subnetworks that they can. Then the problem may not be solvable depending on the value of S, so we introduce ψ to hopefully reduce the number of repeaters. We used distance to discount, assuming people who are far apart are less likely to talk to each other than people which are close to each other. The function we choose for ψ looked like figure 5.

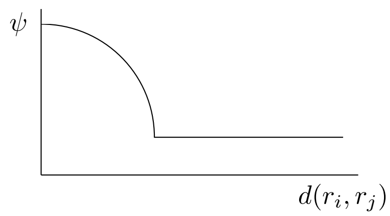

Figure 5: The shape of ψ

3.3.3 Solution

This is a difficult problem to solve because the conditions on the networks depends partly on how the network are chosen as seen in condition 4, so changing one part of the network could have deleterious effects on other parts. We formulated a hill climbing algorithm to compute a solution, with the following steps:

1. First, a subnetwork is created consisting of a single random repeater location and

the value of a heuristic function[1] was computed based on conditions 1, 2, and 3 of the problem.

2. For each subnetwork, the following steps were followed:
   (a) Find all possible repeater locations which could be added to that subnetwork and maintain its connectivity.
   (b) Re-compute the fitness value for a new repeater allocation formed by adding each possible repeater location to the subnetwork, or subtracting a repeater from the subnetwork.
   (c) If any of the possible repeater locations increased fitness, then the one which increases fitness the most is added to the subnetwork and we go to step 2.
3. A new subnetwork is created by looping through all possible repeater locations and finding the one which if made into a 1-element repeater increases fitness the most. Once created, we go to step 2.
4. If no new subnetworks can be created without decreasing fitness then the algorithm terminates.

We did not have time to fully implement this algorithm computationally,

4 The complete model

This section presents the complete problem statement and model.

4.1 Complete problem statement

The problem is the following: Let $D = [0, L_x] \times [0, L_y]$ be a set. Let $P : D \to \mathbb{R}$, $T : D \to \{0, 1, 2\}$, $\psi : D \times D \to [0, 1]$ and $d : D \to \mathbb{R}_+$ be functions. Then the problem is to find a set of repeater locations $R = \{r_1, \ldots, r_N\}$ with a corresponding set of connections $e = \{(r_i, r_j) : d(r_i, r_j) < K\}$ and a collection $\mathcal{R}$ of subgraphs of (R, e) such that

1. Given any point $a \in D$ such that $T(a) \neq 0$ there exists some $r_i \in R$ such that $d(a, r_i) < K$, K some number.
2. $T(r_i) \neq 0, 1$ for all $r_i \in R$.
3. (R, e) is connected and every element of $\mathcal{R}$ is connected.
4. Given any two r_i, r_j, if $\mathcal{R}'$ is the set of elements of $\mathcal{R}$ which connect r_i to r_j, R' is the set of elements of R which are vertices of any element in $\mathcal{R}'$, and $\varphi(r_i, \mathcal{R}')$ is the number of elements of $\mathcal{R}'$ which r_i is a vertex of, then

$$|\mathcal{R}'| \geq \min\{p(r_i), p(r_j)\} + \left(\frac{1}{2} \sum_{r_k \in R', r_k \neq r_i, r_j} \min\{p(r_k), \varphi(r_k, \mathcal{R}')\} \right) \psi(r_i, r_j).$$

[1] Specifically, the heuristic function was defined to be the number of pairs of points which satisfied conditions 1, 2, 3 divided by the total number of pairs of points.

5. If $\mathcal{C} = \{C_1, \ldots, C_{N'}\}$ is some other set of the subsets of R with these two properties then
$$|C_1| + \ldots + |C_{N'}| \geq |R_1| + \ldots + |R_N|.$$

6. The vertex r_i shows up at most S times in the set of subgraphs, where S is some fixed number.

4.2 Complete solution

This is a combination of the cases examined so far and our solution is similar as well. First, we apply our solution to the geography problem to find the set R. Then, we take that set and use it as an input in solving the population problem.

5 Experimental results

We gathered population, elevation, and geographical information for the Puget Sound area in Washington state. We chose this area because of its relative complexity of water and land. Figure 6 shows the different data sets we used.

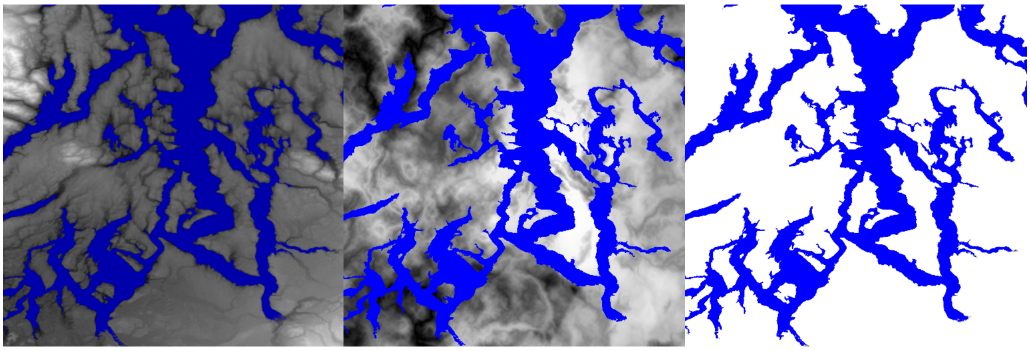

(a) Elevation (b) Population (c) Water/Land

Figure 6: The data used to input into the model

Now we applied the hill climbing solution to the geographical information, as shown in figure 7a. This gave us an upper bound of 60 repeaters necessary to cover the area. Using this number, we applied our genetic algorithm to refine and improve the distribution. The results are seen in figure 7b.

We compared the results of the genetic algorithm to a control algorithm which simply distributed the repeater locations in a uniform grid. The genetic algorithm covered 28% more land than the control algorithm and was much more connected.

We did not have time to compute the hill climbing algorithm and apply it to the problem.

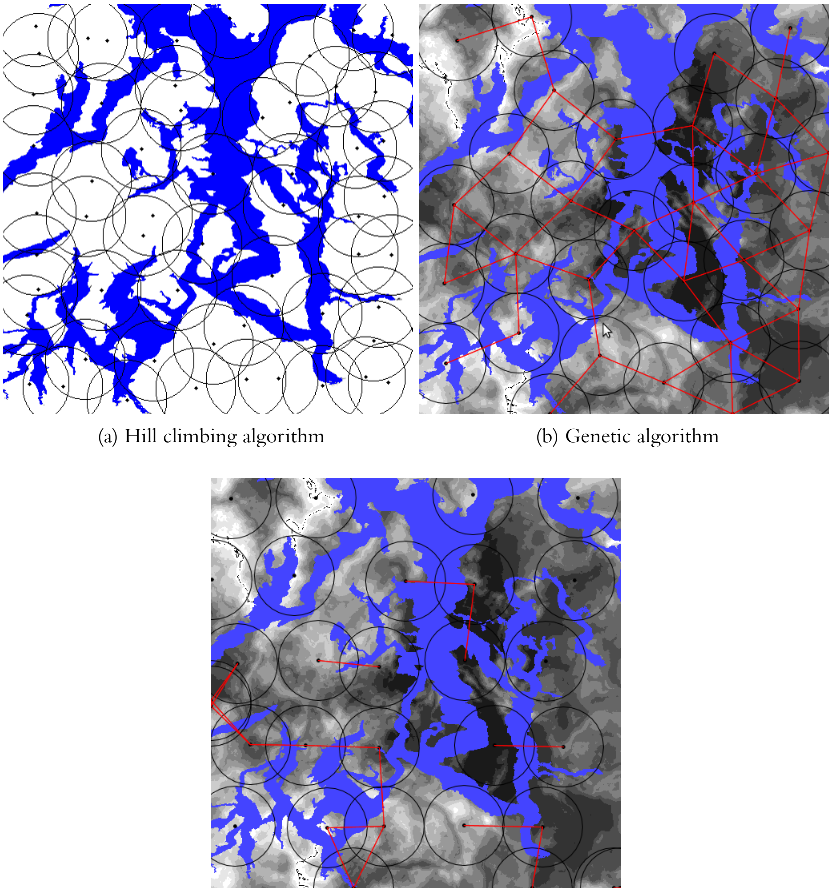

(a) Hill climbing algorithm (b) Genetic algorithm

(c) Control algorithm of equally spaced repeaters

Figure 7: The steps in covering the entire area

6 Further thoughts

Our model did well in distributing the initial network of repeater locations. We did not have time to compute the allocations of subnetworks, however, computation by hand on simpler networks yielded good results. Some areas which we would like to improve and finish are

- Implement the solution to the population problem computationally.
- Experiment with computing the initial repeater allocation using values of d provided by SPLAT!.

A Building repeater networks

In our model we ignored the actual mechanism by which repeater networks are formed. This section elaborates a bit more on why this assumption is reasonable.

Radio repeaters vary widely in their usage design and construction. Many repeaters have capabilities to connect to telephone lines or even to the Internet. Often they are capable of being programmed to respond to different tones. Many have radios which listen for signals on certain frequencies and then tune the repeater to that particular frequency.

We spent some time thinking about how one would construct repeater networks to minimize interference. For example, it is clearly a bad idea to place two repeaters within range of each other, have one listen at frequency F and broadcast at frequency $F + 600$, and have the other listen at frequency $F + 600$ and broadcast at frequency F. We imagine that this would cause a feedback loop.

However, since radio repeaters do appear to be very customizable, and since we are allowed to construct the network however we want, we would construct it as follows: Given some network which we want to construct, we would choose a sub-audible tone not already in use. Then, we would assign each repeater a set frequency of the 6 or so we have to choose from. Each repeater would then be given a small radio which listens for the frequencies which the repeaters surrounding it might broadcast on. Upon receiving the signal it would begin listening on the appropriate frequency and broadcasting on its own frequency.

Regardless of the feasibility of the methods we used, in real life situations we observe a wide variety of functioning radio repeater networks. For example[2] we see repeaters recieving and broadcasting frequencies which are offset by as little as 200 kHz. We also see a plethora of arrangements of repeaters.

The only restriction that this places on repeater networks is that if we have 6 frequencies to use in a network then 7 repeaters cannot be all connected to each other. However, this rarely occurs.

The issue with this setup is that to communicate a user must `key up' their radio and wait until all of the appropriate repeaters have keyed up. However, we think that this is a problem which will always exist over long distances no matter how the radio repeaters are arranged.

B Repeater lengths versus user lengths

Throughout this paper we have been purposefully vague about the issues of repeater signal strength versus user signal strength. Repeaters have powerful signals and are capable of communicating across long distances. Users have weak signals and must be closer to a repeater to communicate.

Based on our observations of maps of repeater signal strengths, it appears that in general a repeater can communicate with another repeater at about twice the distance as a repeater can communicate with a user. Then if a user can communicate with a repeater

[2] See for example a network in Colorado at http://www.colcon.org/fig/colcon_coverage.gif or http://www.colcon.org/fig/colorado_connection_map.gif.

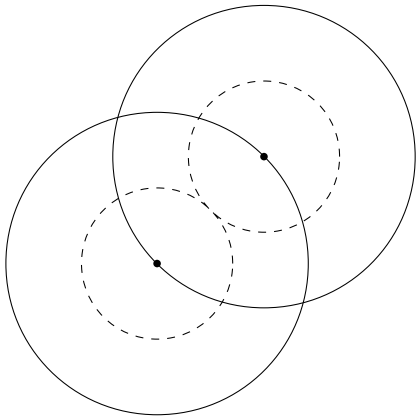

Figure 8: It is okay not to distinguish between repeater strength and user strength

at any point in some area then all of the repeaters in that area must be connected, as demonstrated in figure 8.

The dashed circle represents the radius within which the user can communicate while the solid line represents the radius within which the repeaters can communicate.

References

[1] Minimum-energy broadcast in random-grid ad-hoc networks: approximation and distributed algorithms. *Proceedings of the 11th international symposium on Modeling, analysis and simulation of wireless and mobile systems*. 2008.

[2] Deane, Jason; Terry Rakes and Loren Rees. Efficient heuristics for wireless network tower placement. *Information Technology and Management, V10,N1*

[3] Deming, W. Edwards. Statistical Adjustment of Data. Dover Publications Inc., New York, 1964.

[4] Hufford, The ITS Irregular Terrain Model, version 1.2.2. National Telecommunications and Information Administration, 2002.

[5] SPLAT! Copyright 1997-2011 by John A. Magliacane, `http://www.qsl.net/kd2bd/splat.html`.

[6] Haupt, Randy and Sue Ellen Haupt. Practical Genetic Algorithms. John Wiley & Sons Inc. 1998.

[7] Synman, Jan. Practical Mathematical Optimization. Springer. 2005.

[8] Verblunski, S. On the Least Number of Unit Circles Which Can Cover a Square. *Journal of the London Mathematics Society*, V24 1949.

Repeaters are placed
Subnetworks defined with math
Tomorrow comes sleep