Mobile Network Survivability using Delaunay Triangulations

Jonathan Tse and Eric VanWyk
Olin College

***Abstract*— Delaunay triangulations are used to establish local topology for mobile wireless networks using directional antennae. The maximization of the minimum angle inherent to the Delaunay triangulation is exploited to minimize cross talk, which in turn assures that each individual link can operate at its peak bandwidth. Additionally, the use of mobile relay stations is examined to improve data rate. These mobile relays move themselves into optimal position using information from the Delaunay triangulation and seek to mitigate unavoidably bad node placements that are dictated by the network's users.**

I. Introduction

The advent of several new communication technologies has allowed for a new breed of mobile adaptive networks to become reality. In the current on-demand world, having wireless connectivity to data is paramount. Unfortunately, wireless connections are often poor, due to interference, signal attenuation, or just plain lack of coverage. To make matters worse, the network nodes are often constantly moving. An excellent example is cellphones. Cellphones are always moving because the people that are carrying them are always on the go. Unsurprisingly, cellphone coverage is not uniform; everyone has had ample experience of losing connection just walking down the street. Other communications networks suffer from similar coverage problems.

Cellular relay towers are typically stationary, so dealing with many mobile network nodes that are constantly changing location can be very difficult with stationary towers. These new adaptive networks are designed to "plug" coverage holes by repositioning some relays dynamically, as the other nodes move about. These wireless networks promise to fit user needs better by being easier to deploy and by being able to easily handle change. Of course, wired networks are more reliable, but impractical for these purposes.

A. Networks as Graphs

Any network may be modeled as a graph such that each vertex represents a node—a single device or router, and each edge represents a single method to transfer data between two vertices. In the early telephone network, the vertices were either telephones or switchboards, and the edges were wires suspended by poles. Although the edges in wireless networks have no concrete physical existence, the same graph structure is still applicable.

Using a graph to represent a network enables the use of graph theory to characterize the network. For example, finding a path between two vertices of a graph is analogous to finding a path through the network from node to node. From a network survivability standpoint, finding a cut vertex in the graph immediately identifies high risk points of failure in the network. A cut vertex is essentially any vertex in the graph that, once removed, would turn the graph into two or more disconnected subgraphs. Taking a backbone router offline would have the same effect on any communications network, breaking the network into two or more halves. If the graph has bandwidth capacity information in addition to nodes and their connections, network flow graph theory becomes very useful. Using this kind of graph is a very efficient way to calculate the maximum bandwidth from any node to any other node, assuming there is a path between the nodes.

II. Wireless Networks

Wireless communication has the unique property of signal interference. Since there is no wire to act as a conduit or guide, the signal naturally radiates from the sender. This effect is very similar to the way humans speak. Even though the words (signal) is intended to be heard only in the ears of the listener, everyone else in the area can also hear. This becomes problematic rather quickly when humans congregate. Each send and receive pair must deal with the interference caused by everyone else. If one person begins shouting, it incurs a large cost to those that want to listen to him. As the sum of the costs imposed by other speakers increases, any given conversation becomes less eloquent as it gives way to ”What? Come again?”. Also, as distance increases, wireless signals tend to "spread out" more and thus interfere with other signals, as in Figure 1. The ideal model for wireless transmission depicted in Figure 1 is essentially a cone that widens as distance from the sender increases. Eventually, the distance is great enough that the signal attenuates to the point that it appears as background noise. Also, every jump the signal makes from node to node takes a minimum amount of time, or latency, so the less jumps the signal makes, the better.

Wireless networks suffer all of the problems humans in cocktail conversations, but wireless networks have several additional methods to ensure communication. The most important of these methods is signal routing. Imagine that Alex wishes to send a message to Dylan, but they are on opposite sides of the room. One method to achieve this would be to use a bullhorn and shout above the crowd. Unfortunately, this incurs high cost to everyone else in the room, and interrupts any person in the sound cone of the bullhorn, as in Case 1 of Figure 1. Additionally, the intended recipient may be too far to receive the message, making the signal no more useful than broadcasting interference.

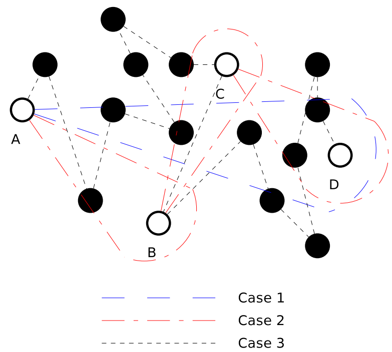

Fig. 1. Each vertex is a person. Vertex A is Alex, vertex B is Bob, C is Charlie, and D is Dylan. All other vertices are other people in the same geographical location at the same time. Case 1 represents the bullhorn approach, where Alex points a bullhorn in Dylan's direction. Case 2 represents the limited-hop approach, where the message gets to the recipient through some number of hops that is less than the total number of vertices in the graph, and Case 3 represents one of the many possible Hamiltonian paths from Alex to Dylan.

A more polite method would be for Alex to ask Bob to ask Charlie to give Dylan his message, as in Case 2 of Figure 1. Since Alex and Bob are within speaking distance, their conversation only impacts a few people around them. This method also accounts for the situation where Alex and Dylan are too far away from Alex to use a bullhorn. By transmitting to another person or node closer to the intended recipient, the sender actually has a chance of getting the signal to the recipient. However, extending this method to the extreme isn't necessarily a good idea either, as in Case 3 of Figure 1. Case 3 represents one possible Hamiltonian path[1] through the graph. Since every jump from node to node increases latency, having the signal bounce through every node incurs a large time cost.

In order to address some of these problems, antenna design is an important aspect of wireless communications.

A. Omnidirectional Antennae

Omnidirectional antennae are probably the most commonplace and cheapest to produce. Continuing the conversation analogy, omnidirectional antennae are akin to shouting loudly in all directions at once. People nearby can hear you just fine, whether they want to or not, but people far away have a hard time hearing you, since your volume has attenuated with distance.

By broadcasting at a given intensity I, a node generates a signal that is received with intensity $I(\theta, r) = Ir^{-2}$. For the sake of modeling, this is a circle of radius I. Clustering many omnidirectional antennae results in a correspondingly high amount of interference; if there are many people close to one another shouting, no one can hear anyone. Geometrically, this corresponds to the areas of each broadcast circle overlapping and causing interference.

B. Directional Antennae

Directional antennae, while less commonplace, are quickly beginning to gain ground due to the unique advantages they provide. A directional antenna has the capability to broadcast a signal selectively in any direction with minimal diffusion. In particular, phased array antennae are particularly good for this because they can change broadcast direction without changing their orientation[2]. The conversation analogy for this is the cupped hands around the mouth while shouting. The cupped hands focus the sound in a particular direction, thus avoiding diffusing the signal. Of course, the signal angle is much better in the electromagnetic realm than the auditory realm.

Directional antenna have notoriously complex emission patterns. They generally consist of one major lobe in the direction of the intended receiver, and many smaller stray lobes. The magnitude—power content—of these stray lobes is low enough such that we may disregard them for this paper and assume ideal transmission. We model this as a cone with dispersion angle $\frac{1}{2}\theta$ and a length of I. Directional antennae are less likely to cause interference than omnidirectional antennae by nature. However, if two or more directional antennae are broadcasting towards the same location, there is the possibility of interference. In this case, the geometric analogy is the overlapping broadcast cones.

C. Maximizing Connectivity and Minimizing Interference

The goal now becomes to simultaneously ensure good network connectivity while avoiding unnecessary interference. The problem can thus be reduced to a question of which edges should be used to populate the graph, and one such way is using planar triangulations.

III. Delaunay Triangulation

A planar triangulation of a set of vertices V can be created by adding edges such that the resulting simple graph G contains $V's$ convex hull and whose interior faces are all triangles. This definition implies that G is maximally planar: no edge may be added without destroying the planarity, and no face contains a vertex. Each set V has many graphs that satisfy these requirements.

Planar triangulations are topologically "bland"[3]. In terms of network survivability, graphs that are "bland" generally remain connected when vertices are removed, since "bland" graphs have a uniform distribution of density[3]. In other words, no part of a "bland" graph is more interconnected than any other part.

[1]A Hamiltonian path in a graph is a path that touches every vertex of the graph only once.

[2]A software algorithm changes the phase of the signal being broadcast by each pole of the phased array antenna. By varying the phase, the signals can interfere constructively or destructively with one another, resulting in a signal that is stronger in only one direction.

[3]The density of a graph is the ratio of the number of edges to the number of vertices.

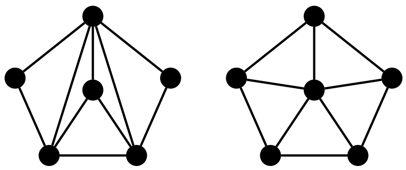

Fig. 2. Two Triangulations of W_5

A triangulation in a 2D plane is considered to be a Delaunay triangulation if the circumcircle of each triangle in the triangulation contains no other points of the graph[2]. In other words, if V is the set of vertices of a graph in a 2D plane and $T(V)$ represents an arbitrary triangulation of V, the circumcircle through the endpoints of any triangle $\nu_1\nu_2\nu_3 \in T(V)$ where $\nu_1, \nu_2, \nu_3 \in V$ contains no other vertices in V.

Figure 2 shows two possible triangulations of W_5. The left triangulation is clearly not a Delaunay triangulation, because there are several triangles whose circumcircles enclose other vertices. The right triangulation is a Delaunay triangulation, since all the circumcircles are empty.

A. Finding the Delaunay Triangulation

One way of finding the Delaunay triangulation is the incremental method, which is a naive algorithm. This method starts by forming a triangle between 3 vertices within V. Additional vertices are incrementally added, and the graph is updated after each addition. All existing triangles are checked and removed if their circumcircle contains the new vertex. The resulting polygonal face is re-triangulated, and the algorithm proceeds to the next vertex, as seen in Figure 3.

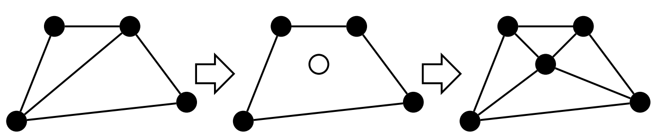

Fig. 3. The incremental algorithm starts with four vertices. An additional vertex is added, and the appropriate edges are removed. Finally, a new triangulation is formed.

It is of interest to note that $DT(V)$is unique only if V is in general position. That is, it is assumed that each set of three vertices in V circumscribe a unique circle[4], and no three vertices in V are collinear. For example, assume that $V_{1\ldots4}$describe the corners of a square. Each set of three vertices chosen from the four describe an identical circle. Their are therefore two Delaunay Triangulations for the graph, as shown in Figure 4. This result bears little consequence on the network topologies, as either choice is equally good. Therefore, we will allow the algorithm to simply choose one if it encounters such a situation.

[4]No four vertices lie on the same circle.

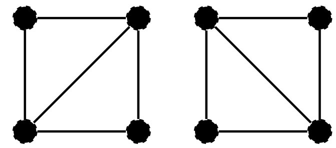

Fig. 4. Two different Delaunay triangulations of a square. Since the vertices of a square are not in general position, the Delaunay triangulation is not unique.

Several more advanced methods may be used to reduce the computation time from $O(n^2)$to $O(n \log n)$, and using a divide and conquer algorithm can achieve $O\left(n \log\left(\log(n)\right)\right)$expected time. However, the basic concept always reduces to finding triangles such that their circumcircles do not contain others.

B. Delaunay Triangulation Benefits

Delaunay Triangulations have several attributes that make them especially salient at this juncture:

1) Each node in V is guaranteed to be connected to its nearest neighbor.[4] Nearest Neighbor links are the least costly to use by most metrics.
2) The Delaunay Triangulation is guaranteed to maximize the minimum angle in the graph[5]. That is, there exists no different triangulation $T(V)$ whose smallest angle is larger than that of the smallest of $DT(V)$. This implies that $DT(V)$avoids creating thin sliver triangles.
3) The minimum spanning tree of the full graph induced by V is always a subgraph of $DT(V)$.[4]
4) As this is a planar triangulation, the number of triangulations and therefore links grow with $O(n)$.
5) The incremental algorithm does not require that all nodes be known at the start of the calculations.
6) Insertions and deletions only have local effects on the triangulation.

C. Mobile Wireless Networks

The benefits detailed in Section III-B make Delaunay triangulations a decent and easy approach to determining a network topology on the fly. As a node is activated by its user, it connects to its nearest neighbor and pushes an announcement of its arrival to the local nodes. The incremental algorithm checks which existing links must be broken, and then reforms links to the new node. The network is protected against the sudden departure of nodes, as it is categorically bland and therefore contains several paths for any data to be routed on. When this condition is detected, only nodes local to the deletion need to establish new links.

IV. Locally Equiangular

Recall the interference pattern of directional antenna in Section II-B. The cone of the signal causes much less interference than its omnidirectional counterpart. However, this is only a start, as directional antennas have a directedness of $2\theta \approx 30^\circ$. By selecting targets that are separated by angle as well as pure distance, it is possible to further reduce the effects of interference. If the minimum angle in any triangle in a triangulation is larger than the directional signal width, it is possible to communicate without creating any effective interference. Note that each transmission will still affect an area, but it is no longer relevant because no nodes are caught within it.

Take a triangle ABC such that $||AB|| < ||BC||$ and $\angle BAC = \phi$. A communicates using a directed antenna with θdispersion on either side of the intended beam. If and only if $\phi < \theta$,B is subject to all communication from A to C.

This creates an additional design constraint concerning these minimum angles. We may choose to interpret this finding in one of two ways. We may maximize the average angle to create better conditions on average. Alternatively, we may maximize the minimum angle to ensure that there is a common standard. The first method is applicable in cases with good data coding that can scale according to SNR[5]. The second method is applicable in cases where there is a threshold effect, in that communication is possible if and only if the SNR is above some value set by the hardware.

For the purposes of this paper "locally equiangular" describes a triangulation that has maximized the minimum angle.

A. Theorem and Proof

Theorem 1: Given a finite set of vertices V and a triangulation $T(V)$, $T(V)$ is locally equiangular if $T(V)$ is the Delaunay triangulation[5].

Proof: Define a triangle ABC. Define another point D to be either inside, outside, or on the circumcircle of the triangle ABC. If the point D is on the circumcircle, we know from Section III-A and Figure 4 that either diagonal of the quadrilateral $ABCD$ will give us a Delaunay triangulation.

If the point D is outside the circumcircle, $\overline{BC}$ will be the diagonal that creates the Delaunay triangulation of $ABCD$, because by definition D lies outside the circumcircle of ABC. The resulting quadrilateral with a diagonal satisfies the locally equiangular constraint detailed in [5], because $\angle BAD < \angle BCD$, $\angle ADC < \angle ABC$, and $\angle DAC < \angle CBD$.

The last case, when D is inside the circumcircle, makes it impossible for ABC to be a triangle in a Delaunay triangulation, because the fact that D is inside the circumcircle of the triangle ABD violates the definition of a Delaunay triangulation. Thus, the other diagonal must be chosen to ensure that no points are contained within circumcircles of triangles. Again, this new quadrilateral satisfies the constraints for being locally equiangular because $\angle BAD > \angle BCD$, $\angle ADC > \angle ABC$, and $\angle DAC > \angle CBD$.

As we can see, the Delaunay triangulation satisfies the criteria for being locally equiangular. As it turns out, the Delaunay triangulation is the only triangulation to satisfy this criteria, as proved in [5]. ■

[5]Signal to Noise Ratio

B. Locally Equiangular Benefits

From Section IV, locally equiangular means to have maximized the minimum angle. Theorem 1 along with its more rigorous proof in [5] proves that the Delaunay triangulation is the de facto way of maximizing the minimum angle, making it the best way of creating a network topology using directional antennae, assuming the cost metric is minimizing potential interference from neighboring nodes in the network.

V. Conclusions on Network Survivability

Recall in Section I that the intended application is mobile network survivability. Section IV shows that Delaunay triangulations minimize interference when using directed antennae with all the benefits of a planar triangulation. Clearly, the locally equiangular benefit is a major one, because ensuring signal integrity ensures that links—edges in graph theory terms—remain active and can communicate at close to their maximum bandwidth.

Using the incremental algorithm described in Section III-A, adding new nodes or moving nodes around in time only affects the local subgraph. This is particularly beneficial in large scale network deployment, because small changes do not affect the overall network integrity, i.e. the addition, removal, or repositioning of nodes affects only a handful of nodes nearby. This phenomenon is directly related to the fact that planar triangular graphs are bland—the network can survive single node changes long enough for the network to re-triangulate an optimal solution.

Also recall that the mobile network described earlier involves mobile relay towers. Choosing the location to add new nodes—vertices—can also be accomplished with Delaunay triangulations, as described in [2]. The center of the largest circumcircle of a Delaunay triangulation is an ideal candidate for placing a relay tower, since the largest circumcircle touches nodes that are the farthest apart. Since the nature of wireless signals is that they attenuate proportional to distance, putting a relay in between the largest distances shortens the distance from node to node, thereby increasing the chance that a signal will actually get through[2].

References

[1] M. Bern, D. Eppstein, and F. Yao. "The expected extremes in a Delaunay triangulation," in *International Journal of Computational Geometry and Applications*., 1(1):79–91, 1991. http://citeseer.ist.psu.edu/bern91expected.html

[2] C. Wu, K. Lee, and Y. Chung. "A Delaunay triangulation-based method for wireless sensor network deployment," in *Proceedings of the* 12^{th} *International Conference on Parallel and Distributed Systems*, 2006.

[3] John G. Michaels & Kenneth H. Rosen. *Applications of Discrete Mathematics*. McGraw-Hill, Inc. New York, 1991.

[4] R.Press. "Lecture 17" in Lectures, Computational Geometry, 2003. http://www.cs.wustl.edu/~pless/506/l17.html

[5] R. Sibson. "Locally equiangular triangulations," in Computer Journal. 21:243-245, 1978.

Appendix A: Notes

R. Sibson's paper, "Locally equiangular triangulations," proved to be impossible to find, and heavily referenced, making writing this paper difficult. The paper ended up costing $23.00, but we didn't make the decision to buy it until late into the writing process. We have included the paper for your convenience.