Optimization of Directional Antenna Network Topology in Airborne Networks

G. Hadynski, S. B. Lee, G. Rajappan, R. Sundaram, X. Wang, F. Zhou

Abstract—

Future IP-based Airborne Networks, important components in net-centric military communications, are envisioned to consist of a persistent backbone core network and dynamic tactical edge networks. The backbone would consist of quasi-stable platforms equipped with multiple high-capacity directional wireless links. The tactical edge networks would consist of highly dynamic platforms such as fighter jets equipped with omni-directional wireless links, and these would be interconnected by the backbone core network. Maintaining optimal backbone topology is an important problem with significant operational impact. Factors such as non-uniform link capacities, the number of traffic sources and sinks, and connectivity complicate the problem. The solution consists of making optimal selection of the link directionality and the possible insertion of communication relay nodes.

We approach the solution by abstracting the network as a template from which to select the optimal combination of edges (transmitter-receiver pairs) and nodes (relays). Through innovative graph and flow-theoretic reductions we show that the single sink (or alternatively single source) case can be solved in polynomial time for uniform backbone link capacities. In contrast, we prove not only that the problem is NP-complete for non-uniform backbone link capacities but that the non-uniform case of the problem is hard to approximate to within even a logarithmic factor. Nevertheless we present a scheme based on iterative rounding that scales well in practice. Simulations demonstrate that our algorithm achieves a performance within a factor 2 of the theoretical best. This allows us to conclude that the use of algorithmic techniques in configuring backbone networks can contribute significantly in improving network performance.

Gregory Hadynski (email: Gregory.hadynski@rl.af.mil) is with Air Force Research Laboratories, Rome NY USA.

Fangfei Zhou (email: youyou@ccs.neu.edu) and Ravi Sundaram (email: koods@ccs.neu.edu) are with the College of Computer Science, Northeastern University, Boston, MA USA.

Seoung Bum Lee (email: sblee@mayflowercom.com), Gowri Rajappan (ph: 781-359-9500 email: rajappan@mayflowercom.com) and Xiaofei Wang (email: wang@mayflowercom.com) are with Mayflower Communications Company, Inc., Burlington, MA USA.

The work presented herein was performed under the US Air Force Research Laboratory Contract No. FA8750-09-C-0119.

Approved for Public Release; Distribution Unlimited: 88ABW-2010-2077 dated 16 April 2010.

Index Terms— **Heterogeneous networks; Mobile ad hoc networks; Network architectures; Network optimization; Network planning and topology.**

I. Motivation

There is immense ongoing interest in the development of an IP-based Airborne Network based on mobile ad hoc networking concepts [16]. Such a network may consist of disparate mobile airborne platforms; will provide reliable ad hoc interconnectivity to terrestrial and airborne nodes; and will be applicable to net-centric military communications. A wireless network that includes an airborne component is architecturally reminiscent of conventional wireless telecommunications networks. Relatively quasi-stable airborne platforms, equipped with directional data links, could form a high capacity backbone core akin to the backhaul in telecommunications networks. This directional backbone core would serve as a reliable data transport infrastructure for highly dynamic terrestrial and airborne platforms at the network edge, which are typically equipped with lower capacity omni-directional radios. The latter are akin to end users in conventional telecommunications networks.

As in any mobile network, the performance of such a network depends to a great extent on the prevailing network topology. The network topology determines the availability and the characteristics of the data path between any two nodes of the network. These characteristics directly dictate the performance of data transmission, such as throughput and delay, between the nodes. For a network of fixed nodes the topology is essentially static. It is determined by deployed physical connections between nodes. For a network with mobile nodes the topology is dynamic with several topology choices possible. Methodologies should be devised for active topology management and optimization in the Airborne Network with the aim to improve overall network performance.

In this paper, we are interested in the management and optimization of the backbone core of the Airborne Network. This is an important problem because the backbone core network, conceived as the bedrock on which the Airborne Network is constructed, is perhaps the biggest determinant of the network performance. A well-optimized backbone core that maximizes capacity will allow the network to scale and will provide highly assured transport to a large number of concurrent traffic flows. In contrast, a poorly constructed backbone core will limit the network scalability and may introduce unacceptable packet losses and latencies at even nominal traffic loads.

Optimization of the directional backbone core is a challenging proposition due, in particular, to the following factors: (1) The nodes have a fixed number of directional links (and this number may be different for different nodes), (2) The links may be of different data rates (for instance, due to differences in received Signal to Interference and Noise Ratio (SINR)), (3) The incoming and outgoing links may be of different radio types (in which case the node is assumed to be able to translate between the corresponding waveforms), and (4) The link bandwidth may be shared between bidirectional traffic by means of, for instance, a time division multiplexing protocol. The optimal topology solution, given one or more of the above constraints/factors, consists of the link configurations that optimize network performance in terms of capacity and latency. (Additional performance considerations such as survivability are of interest as well, but discussion of such measures will be deferred to follow-on publications.) By link configurations we mean the link directionality (set by appropriately orienting the antenna beam) and the insertion of additional links (accomplished by strategic insertion of relay nodes).

The inherent complexity of the directional radio network topology problem necessitates rich models that can capture the variety of constraints described above as well as the complete solution space. We propose an inherently rich system model based on graphs that effectively capture the backbone core of the Airborne Network. Using this model, we are able to characterize the topology optimization problem and form conclusions about the optimal topology solution. We are also able to employ flow techniques to obtain good solutions, which approach the optimal solutions at a fraction of the complexity. Since these low-complexity topology solutions can be implemented in real time for large networks, we enable an Airborne Network that continues to provide excellent performance in spite of network dynamics and growth.

In Section II, we summarize the original contributions of this work. In Section III, we survey related work. In Section IV, we present our model, its formulation, and the overall terminology. In Section V, we describe graph gadgets and compositions that is at the core of our model. In Section VI, we examine the special case of a directional network with uniform capacity links. In Section VII, we examine the general case of a directional network with non-uniform capacity links. In Section VIII, we present a low-complexity heuristic for topology solutions that approach the optimal solution and present simulation results that showcase its performance. In Section IX, we summarize conclusions and future work.

II. Our Contributions

We present a practically useful and mathematically well-defined formulation of the problem of optimally configuring a backbone network. We characterize the computational complexity of the problem and some special cases through algorithms as well as hardness results. We present an efficient heuristic for the general case and demonstrate its usefulness through simulation. At a high level we present a collection of *graph gadgets* along with techniques for combining them that enable the problem to be tackled using well known techniques from the theory of flows on networks. A detailed summary of our major contributions is:

- We consider a general situation with nodes of different processing capacities, possessing radios of different types and ranges, links that can be uni-directional or bi-directional and relay nodes that can be inserted at pre-specified locations. We show how the general problem can be precisely modeled by composing various graph gadgets appropriately. The resultant problem is an example of an MIP (Mixed Integer Programming) problem.
- We show that the special case problem where all capacities are uniform can be solved by using well-known algorithms for minimum cost maximum flow.
- We show that the general (nonuniform) case of the problem with just *two* different link capacities is NP-complete. We amplify our (NP-completeness) construction and show that the problem is in fact inapproximable to within a logarithmic factor.
- We present a practical scheme based on iterative rounding for the general problem. We experimentally evaluate our scheme on synthetic topologies with capacities and constraints drawn from practical systems. We show that our scheme is both performant and scalable.

Our modeling techniques are relatively straightforward and we do not claim novelty for the graph gadgets themselves so much as the idea of applying them to capture the rich variety of situations involving wireless backbone networks using directional antennae and relays.

III. Related Work

Network design in general and wireless networks in particular have been studied extensively [14]. The general approach of modeling networks as graphs and solving the routing problem using flows has a vast literature [1]. The theory of flows first developed in [7] has subsequently been generalized by developments in the theory of Linear Programming [13]. Though well-known polynomial-time algorithms such as the Ellipsoid Algorithm or Interior-Point Algorithms [10][15] are known, nevertheless it is still an open problem to find a polynomial-time *combinatorial* algorithm for Linear Programming [9]. The algorithms presented in this paper are polynomial-time combinatorial algorithms. [11] is one of the most efficient implementations of a minimum cost maximum flow in the public domain. We utilize this implementation in our simulations. [9] is the standard reference on the theory of NP-completeness and our hardness results follow in this tradition. Our NP-completeness and inapproximability results use arguments similar to those used in the area of degree bounded network flows [3][6]. The area of topology management in ad hoc networks has received much attention over the past decade. [2] has focused on the problem of energy conserving topologies using minimum dominating sets as abstractions of high traffic gateway nodes. [10] has focused more specifically on the problem of discovery and topology management in wireless networks with directional antenna. [10] studies the problem from the viewpoint of minimizing the stretch or hop count versus the localized degree of nodes. Our focus in this paper is not on discovery. We assume that nodes are aware of their neighbors and their radio capabilities. We also assume that the information is centrally available and focus on finding the optimal configuration of the network, subject to a variety of radio and range constraints,that maximizes the concurrent bandwidth from the different sources to a single sink.

IV. Model, Formulation and Terminology

We model a network as a directed graph G = (V, E) with n = |V| vertices (nodes) and m = |E| directed arcs (links). We use lowercase letters to denote the individual elements, i.e. V = $\{v_1, v_2, \ldots v_n\}$ and E = $\{e_1, e_2, \ldots e_m\}$, where each $e_k = (i, j)$ is the directed arc from v_i to v_j. Arcs have capacities and costs. We use r_{ij} to denote the rate or capacity, and c_{ij} to denote the cost of arc $e_k = (i, j)$.

The practical problem we wish to tackle is as follows: we are given a wireless network consisting of nodes with radios of different types and directional antennas; we have to configure the network so as to achieve the highest *concurrent* bandwidth from a specified set of source nodes to a sink node. This accurately captures a variety of real-world situations where a collection of sources is generating information (e.g. video feeds from Unmanned Aerial Vehicles) and the information needs to be aggregated at a single sink (Command Center). For the purposes of this work we assume all the sources have the same demand and we focus on simultaneously maximizing the largest fraction of each demand, i.e. maximizing concurrent bandwidth. The problem of maximizing concurrent bandwidth from a given set of sources to a sink is easily captured as a linear program [12][14] and solvable in polynomial-time. However the problem rapidly gains in complexity when we consider that the nodes are spatially organized with radios of differing compatibilities and ranges. Because each node has a fixed number of radios of certain types it can communicate only with certain other nodes. Further, if a node uses a certain radio to communicate with one node then it can no longer use that same radio to communicate with a different node because the assumption here is that the radios use fixed directional antenna. These constraints sufficiently complicate the problem so that it is no longer a matter of simply figuring out how to route the maximal flows but more one of how to configure the links in the first place so as to enable the maximum concurrent bandwidth.

V. Graph Gadgets and Compositions

We now consider the different kinds of constraints imposed by a wireless network with nodes having radios of different types. We also consider processing constraints at the nodes themselves as well. Lastly, we consider a situation where we have a fixed number of relay nodes that we wish to deploy so as to maximize the concurrent bandwidth from the set of sources to the sink. We now discuss each of the constraints and present graph gadgets that model them.

A. Radio constraints

The Airborne Network would typically include different radio types. We assume that there are a fixed number of radio types. A radio of one type may be used only to communicate with a radio of the same type. Radios vary in the communication rates they are capable of sustaining. Each node is assumed to have a collection of radios of different types. Links are based on directional antenna so that if a node uses a radio of a certain type to communicate with another node then it may not use that same radio to communicate with a different node. We assume that the nodes are located in space and naturally there are range constraints on which nodes may communicate with another. We model the entire set of possibilities as a *template* graph. In other words, the graph represents the entire space of possibilities and we use graph gadgets to represent the constraints on radio types available at a node. The graph gadget is best explained through an example. Suppose a node has two transmitters of type I and one receiver of type II then we would set up an outgoing arc (termed the *pooling* arc) of capacity 2 to which we connect all the nodes to whom this node can transmit using a type I transmitter and an incoming arc (pooling arc) of capacity 1 to which we connect all the nodes from whom this node can receive using a type II receiver.

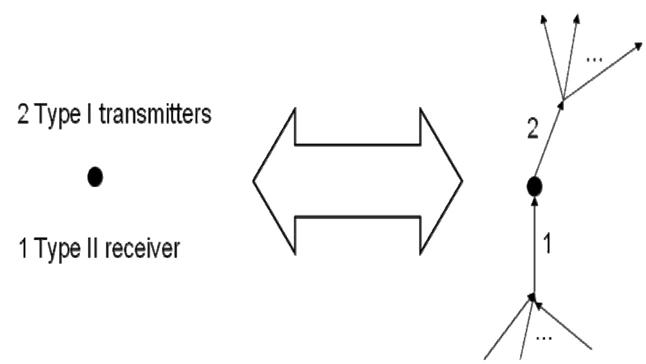

Figure 1: Graph gadget that incorporates the radio restrictions at a node. Arc capacities are shown next to arcs.

Lemma The gadget in Figure 1 captures the radio restrictions at a node.

Proof (Sketch): Even though a node may have any number of possible candidate links it can participate in with a radio of a given type, nevertheless, because of the single arc representing that type in the gadget the total number of radios used cannot exceed the capacity of the arc. In this way the gadget transforms a budget on the number of radios of a given type into an arc capacity constraint. Observe that an omni-directional transmitter can be modeled by an outgoing arc of infinite capacity. ■

B. Node-capacity constraints

Consider situations in which nodes have constraints on the amount of data they can handle. Such situations could arise for example when the node has to translate between different radio interfaces or filter the data passing through it for security reasons or archive the data. We create a graph gadget for each node with a node-capacity constraint. Again the gadget is best explained through an example. Let us assume that 1 unit represents the transmission capacity of the link with the lowest capacity in the network. Then if a node can handle at most 10 times this capacity then we represent(replace) this node by two nodes, a head and a tail with an arc from the tail to the head of capacity 10; we connect all incoming arcs to the tail and all outgoing arcs to the head.

Lemma The gadget in Figure 2 captures node capacity constraints.

Proof (Sketch): Observe that the intermediate arc from the tail to the head aggregates all the incoming flow and channels it to the outgoing arcs. Hence, the aggregate flow through this arc will never exceed the node capacity constraint. ■

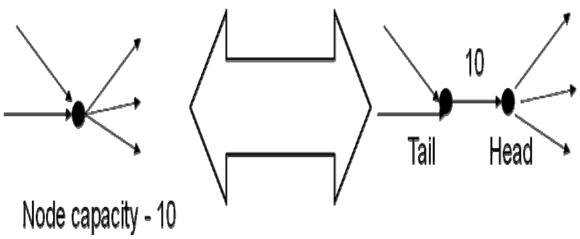

Figure 2: Transformation from node with capacity restriction on left to graph gadget on right. Node is split into a tail and a head with a capacity restriction on the intermediate arc.

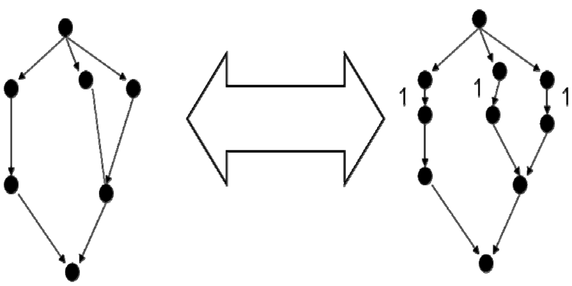

Figure 3: Instance on left has 3 candidate locations (the 3 nodes just below the source) but only 2 relays. Transformation on right captures relay restriction with budget of 2 on cost.

C. Relay constraints

In a general Airborne Network there are a certain number of relay nodes or platforms that can be inserted into the network to improve the performance of the network. Typically, the number of candidate locations where the nodes may be inserted will exceed the number of relay nodes available and so the problem becomes one of figuring out the optimal subset of locations for insertion. In the general case the relay nodes may differ in their capacities as well as the radio types they possess; there may also be different costs of insertion. In this paper we restrict our attention to the situation where all relays are identical and have the same insertion cost. We express the relay constraint using a graph gadget that uses arc costs. As with radio constraints and node-capacity constraints, relay constraints are best explained with an example. Suppose we have 2 relays in total and there are 3 possible locations where they can be utilized. Then we replace the node at each of those 3 places in the template graph by two nodes, a tail and a head, with an arc from the tail to the head with cost 1; we set the total budget for a minimum cost maximum flow to be 2.

Lemma The gadget in Figure 3 transforms the budget on relay nodes to a budget on the minimum cost maximum flow.

Proof (Sketch): Observe that in any integral flow the total cost is equal to the number of relays whose intermediate arcs support positive flows. Hence there is a minimum cost flow with respect to the budget iff there is a collection of relay locations where relays can be placed to achieve the flow.■

D. Bi-directional links

We now demonstrate a gadget for capturing bi-directional links with shared capacities. Wireless links using frequency or time division multiplexing share capacity in that the sum total of the capacities in the two directions (uplink and downlink) is a fixed constant. In other words the link can be thought of as two arcs between the nodes v_i and v_j, (v_i, v_j) and (v_j, v_i) but with capacities summing up to a constant, i.e. $r_{ij} + r_{ji} = r$. This is not an uncommon situation and we now show how any such link can be represented using an appropriate graph gadget. Again we use a diagram for our illustration.

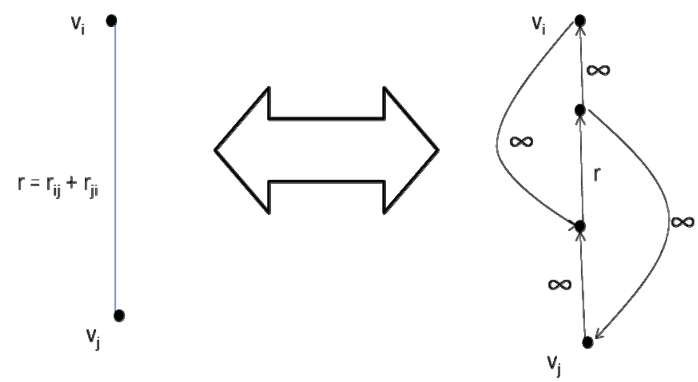

Figure 4: Bi-directional link with shared capacities and its corresponding gadget. Arc capacities are shown next to arcs.

Lemma The gadget presented in Figure 4 exactly captures a bi-directional link with shared capacities.

Proof (Sketch): Observe that the total flow from v_i to v_j and from v_j to v_i in the gadget can never exceed r because it is limited by the vertical straight arc in the middle. ■

VI. UNIFORM CAPACITIES – SPECIAL CASE

We now consider the special case where all the capacities are uniform. By uniform we mean that all link capacities (irrespective of radio type) are the same, say 1 unit. We also assume that relay nodes have a node capacity of 1. It might seem overly restrictive to have all link capacities be the same but it really is not, the model still has a lot of richness to it. We still support different radio types with compatibility constraints, node capacity constraints (which are a multiple of the link capacity), relay constraints and bi-directional links.

Given a backbone network with its attendant constraints we model it using the graph gadgets detailed in the previous section to create the uniform BACKBONE-CONFIG problem. We then solve this minimum cost maximum flow problem on the transformed network and the resulting flow gives us the optimal configuration of the backbone network.

Theorem The BACKBONE-CONFIG problem with uniform link capacities has a solution satisfying all the demands iff there is a maximum flow of concurrent bandwidth 1 unit and a total cost that does not exceed the relay budget.

Proof (Sketch): First, we apply the various graph gadgets from the previous section to the given instance. We then create a super-source and connect it to all the sources with an arc of capacity 1. Observe that all link capacities are equal and all other arc capacities are multiples of the link capacity. According to the theory of flows with *integral* capacities there exists an *integral* minimum cost maximum flow [7]. Hence in any maximum flow every link has either 0 flow or 1 flow. Further, the maximum flow from the super-source is equal to its outdegree iff all of the original demands can be concurrently satisfied. And the minimum cost of such a maximum flow is equal to the number of relay nodes utilized, i.e. relay node locations where the intermediate arc supports a 1 flow. Hence the theorem follows. ■

Corollary The BACKBONE-CONFIG problem with uniform link capacities can be solved in time $O(\min(n^{2/3}, m^{1/2}) * m * \log(n^2/m))$ [12]. ■

VII. NON-UNIFORM CAPACITIES – GENERAL CASE

First, we present an NP-completeness result for the decision version of the problem in the non-uniform case. Amplifying the basic NP-completeness reduction we then present a more involved reduction showing that the search version of the problem cannot be approximated to better than $(1/2)\log_2 n$, where n is the number of nodes.

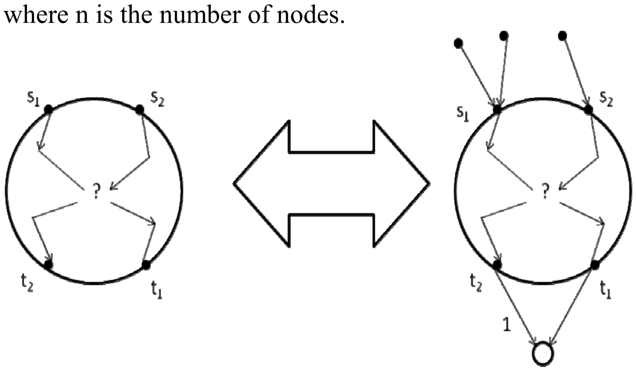

Figure 5: Transformation from 2DIRPATH on left. Top three nodes at right are sources with demand 1. Except for left arc into the sink at bottom of capacity 1 rest have capacity 2. All nodes in the circle have radio constraint of at most one outgoing link.

A. NP-completeness

Theorem The BACKBONE-CONFIG problem with non-uniform link capacities is NP-complete.

Proof: It is easy to see that the problem is in NP because we can guess the collection of links satisfying the radio constraints and then verify using Linear Programming [9] that the set of demands can be simultaneously satisfied.

To show that the problem is NP-hard we show a reduction from the 2DIRPATH problem, which is shown to be NP-hard in [6] and defined as follows: Given an n-node directed graph G and two node pairs s_1, t_1, s_2, t_2 find node-disjoint directed paths from s_1 to t_1 and s_2 to t_2. We reduce this problem instance to an instance of the BACKBONE-CONFIG problem with radio constraints as shown in Figure 5.

It is easy to see that there exists an orientation of the antennae and configuration of the radios satisfying all the demands iff there exist node disjoint directed paths in the original graph. ■

In fact, the above reduction also shows that if all the demands cannot be satisfied then at most a 2/3 fraction of each demand can be satisfied thus giving us the following:

Corollary The search version of the BACKBONE-CONFIG problem with non-uniform capacities cannot be approximated to better than 2/3 factor.

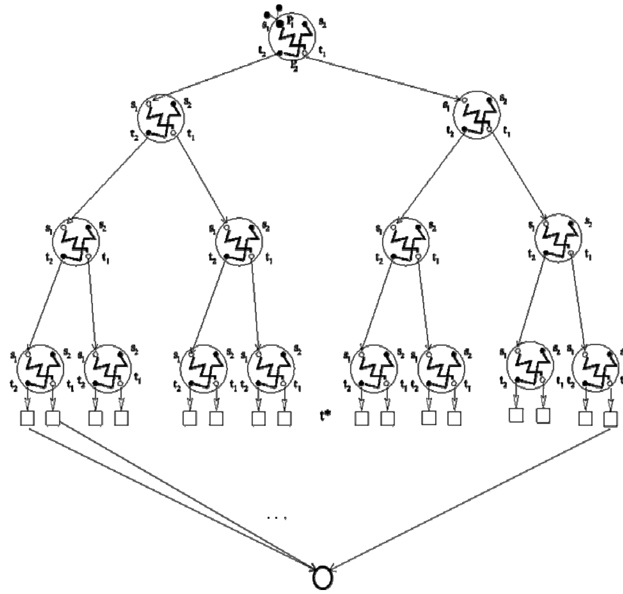

Figure 6: Each node of a binary tree is replaced by an instance of 2DIRPATH. The sink is at the bottom.

B. Logarithmic inapproximability

Theorem The search version of the BACKBONE-CONFIG problem with non-uniform link capacities cannot be approximated to a factor better than $\Omega(\log n)$ on an n-node graph.

Proof(sketch): Due to space constraints we give a brief sketch here. As shown in Figure 6, we take a binary tree with n nodes and replace each node with a copy of the transformed 2DIRPATH from our earlier NP-completeness reduction. As before, we set up the radio constraints so that nodes have at most one outgoing link. We set all links to have capacity 2 except for the two links into s1 at the root which have capacity 1 from two nodes with demands 1. We set all the s2, t2 nodes to have demand 1. From the construction it is easy to see that if there exist node disjoint paths in 2DIRPATH then all the demands are satisfiable. But if there exists no solution to 2DIRPATH then because nodes are constrained to have at most one outgoing link the demands starting from s1 at the root will cumulate all the way to the sink at the bottom creating a congestion of (log2 n)/2 at the bottom which gives us our result.

VIII. HEURISTIC AND EXPERIMENTS

From the previous section we see that the general problem is very hard. Now we present a heuristic, ITER-ROUND based on iterative rounding for the general case. The basic idea of the heuristic is to set up the network using the graph gadgets and then solve for the minimum cost maximum flow (using the code from [11] as a subroutine). The problem however is that the solution may be fractional. We then pick the most heavily loaded link (as a percentage), breaking ties arbitrarily and round up the flows on them. This rounding up affects two sets of arcs – the radio constraint arcs and the relay node arcs. We commit to those arcs that get rounded up, i.e. in the case of relay nodes we reduce their arc costs to 0 and reduce the cost budget for the flow problem by the arc cost, and in the case of transmission links we directly connect the transmitting and receiving nodes while reducing the capacity of the pooling arcs by the capacity of the corresponding link. Note that we always produce a feasible solution in this way, though we may fall short of optimality.

ITER-ROUND

- Set up network using graph gadgets
- Repeat until no flow
  - Compute minimum cost maximum flow
  - Round up most heavily loaded link or relay node
  - Modify graph gadgets to be consistent with rounding
- Output optimum flow from sources to sink

In order to study the performance of ITER-ROUND using simulations, we placed n nodes at random on a square, 10 units by 10 units. We assumed there are two radio types: red and blue. Red (blue) radios can only communicate with red(blue) radios. We gave each node 1 or 2 blue/red transmitters/receivers uniformly at random, i.e., the average node has blue/red out/in degree of 1.5. Two nodes were set to be in communication range if they were within distance 2 of each other. We did not simulate node capacities or bidirectional arcs. We ran a number of simulations for n ranging from 200 to 500 in steps of 50. We ran 10 simulations for each n, for a total of 70 simulations. For each graph we computed the minimum cost maximum concurrent flow using ITER-ROUND.

We also computed the minimum cost maximum concurrent flow using brute force approach of trying all possible link combinations so that we could compare the quality of the solutions produced by ITER-ROUND. We found that on the average ITER-ROUND produced solutions that were at most a factor 2 worse. Of course ITER-ROUND is significantly faster than the approach of trying all possibilities using brute-force. See Figure 7 for a graph of the quality of the solution produced by ITER-ROUND as a function of n.

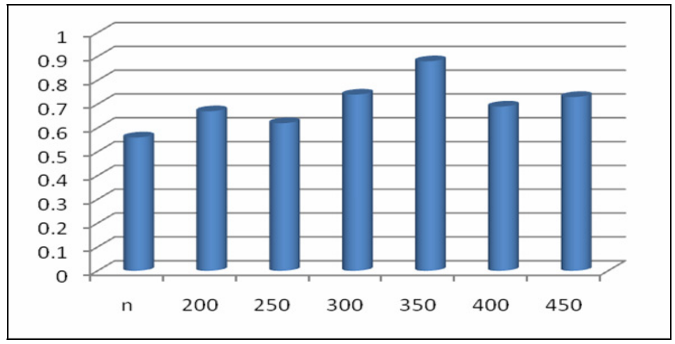

Figure 7: Quality of solution produced by ITER-ROUND represented as a fraction of optimum vs. number of nodes.

IX. Conclusions and Future Work

We have initiated a study of directional antenna based topology management in Airborne Networks. We created a well-defined formulation of the problem and showed that the general case is provably hard. We showed that the uniform case has efficient algorithms. We also presented a heuristic for the general case based on iterative rounding and showed that it has good performance on the average for real-world scenarios. Interesting directions for future research include additional performance considerations such as survivability, efficient distributed algorithms, and the problem of scheduling transmissions taking interference into account.

References

[1] Ahuja, R., Magnanti, T., and Orlin, J., “Network Flows: Theory, Algorithms and Applications”, Prentice-Hall, 1993.
[2] Bao, L., and Garcia-Luna-Aceves, J., “Topology Management in Ad Hoc Networks,” Proceedings of MOBIHOC, pp. 129-140, 2003.
[3] Chen, J., Kleinberg, R., Lovasz, L., Rajaraman, R., Sundaram, R., and Vetta, A., “(Almost) Tight Bounds and Existence Theorems for Confluent Flows,” Journal of the ACM, 54(4), 2007.
[4] Chen, J., Rajaraman, R., and Sundaram, R., “The Confluent Capacity of the Internet: Congestion vs., Dilation,” IEEE ICDCS, 2006.
[5] Chhabra, P., Laoutaris, N., Rodriguez, P., and Sundaram, R., “Algorithms for Constrained Bulk-Transfer of Delay-Tolerant Data,” Proceedings of IEEE ICC, 2010.
[6] Donovan, P., Shepherd, B., Vetta, A., and Wilfong, G., “Degree-constrained Network Flows,” ACM STOC, pp. 681-688, 2007.
[7] Ford, L., and Fulkerson, D., “Flows in Networks,” Princeton University Press, Princeton, NJ 1962.
[8] Fortune, S., Hopcroft, J., and Wyllie, J., “The Directed Subgraph Homeomorphism Problem,” Theoretical Computer Science, 10(2), pp. 111-121, 1980.
[9] Garey, M., and Johnson, D., “Computers and Intractability: A Guide to the Theory of Incompleteness”, Freeman, 1970.
[10] “Gelal, E., Jakllari, G., Krishnamurthy, S., and Young, N., “Topology Management in Directional Antenna-equipped Ad Hoc Networks,” IEEE Transactions on Mobile Computing, 8(5), pp. 590-605, 2009.
[11] Goldberg, A., Network optimization library. Available at http://www.avglab.com/andrew/soft.html
[12] Goldberg, A., Rao, S., “Beyond the flow decomposition barrier,” Proceedings of FOCS, 1997.
[13] Grotschel, M., Lovasz, L., and Schrijver, A., “Geometric Algorithms and Combinatorial Optimization” Springer-Verlag, 1988.
[14] Kurose, J., and Ross, K., Computer Networking: A Top-Down Approach”, Addison-Wesley, 2009.
[15] Schrijver, A., “Theory of Linear and Integer Programming”, Wiley, 1998.
[16] HQ ESC/NI1, “Airborne Network Architecture,” 7 October 2004. Available at http://herbb.hanscom.af.mil/Hot_Buttons/Airborne_Networking/AN_Architecture_7_Oct_2004.doc