VillageNet: A low-cost, 802.11-based mesh network for rural regions

Partha Dutta† Sharad Jaiswal† Debmalya Panigrahi†
K V M Naidu† Rajeev Rastogi† Ajay Todimala§

†Bell Labs Research India
Bangalore, India
{pdebmalya,parthad,jsharad,naidukvm,rastogi}@lucent.com

§Dept. of Computer Science & Engineering
University of Nebraska - Lincoln
ajayt@cse.unl.edu

***Abstract*— VillageNet is a wireless mesh network that aims to provide low-cost broadband Internet access for rural regions. The cost of building the network is kept low by using off-the-shelf IEEE 802.11 equipment and optimizing the network topology to minimize cost. In this paper we describe the over-all operation of VillageNet and discuss two fundamental problems in building such a network.**

Nodes in VillageNet communicate using long-distance point-to-point wireless links that are established using high-gain directional antenna. VillageNet uses the *2P* MAC protocol [17], that is suited for the interference pattern within such a network. However, the 2P protocol requires the underlying mesh graph (for each 802.11 channel) to be bi-partite. Thus, if K channels are available, then an important consideration is how to select K bi-partite subgraphs to activate, such that the demands of the nodes are best met. We formally pose this problem and present some initial results.

Second, we observe that the dominant cost of constructing such a mesh network is the cost of constructing antenna towers at nodes. The cost of a tower depends on its height, which in turn depends on the length of its links, and the physical obstructions along those links. Thus to minimize cost, we pose the problem of deciding which links should be established, such that all villages are connected and the cost of constructing antenna towers to establish the selected links is minimized.

I. Introduction

In this paper, we describe VillageNet, an IEEE 802.11 based long-distance mesh network for rural areas. We outline the architecture of this network (based on the work done in the Digital Gangetic Plains project [2]). We discuss the routing and channel allocation problem and the topology construction problems in the these networks, and present some initial solutions.

Rural networks exist in population areas with very low paying capacity. Hence a very important requirement for these networks is to minimize the infrastructure costs. The cost of laying wire to rural areas is prohibitively expensive and this is also true of traditional wide area wireless technologies such as cellular networks and upcoming technologies like IEEE 802.16 Wi-Max. Recent work [17] in the literature has demonstrated an alternative approach by building rural mesh network prototypes using IEEE 802.11 Wi-Fi equipment. IEEE 802.11 equipment is highly commoditized and goes a long way in addressing the cost issues for this environment, and is thus the cheapest option to build such networks.

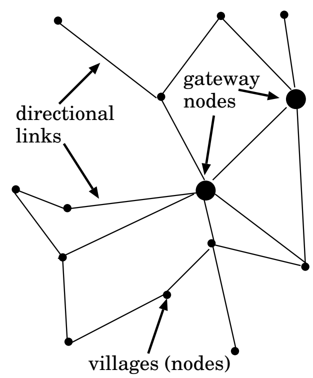

Fig. 1. A wireless rural mesh network

A. Architecture

A typical rural mesh network is depicted in Figure 1. It would consist of a cluster of villages connected with each other through point-to-point wireless links. Some special nodes in this mesh, called *gateway* nodes, will be connected to the wired internet. Other mesh nodes will connect to the gateway node (and thus, to the rest of the internet) through one or more hops in the mesh. Rural mesh networks are characterized by a fixed, outdoor topology (a node in this network will be a village) and very long-distance links between the nodes (about 10-15kms).

The long-distance point-to-point links are established using IEEE 802.11 equipment, and high-gain ($\tilde{2}$4dbi) parabolic grid directional antennas. The 802.11 MAC was originally designed for (and widely deployed in) short-distance campus area networks with mobile nodes, and thus may not be well suited for deployment in the long distance. But, as [17] has demonstrated, by using high-gain directional antennae, a line-of-sight long-distance ($>$ 10 kms.) 802.11 link can be established.

B. Interference Model

We described above that nodes in our mesh networks will communicate with each other using directional antennae. While directional antennae are designed to transmit and receive in a specific direction, the directionality of this radiation becomes effective only at longer distances from the sender. This is also called the *near field effect*. Because of this effect, adjacent links at a node interfere with each other in certain communication modes. Specifically, at any node, simultaneous transmissions and receptions on the *same channel* are not possible − the transmissions will interfere with the receptions. This is called *Mix-Rx-Tx interference* [17].

Apart from Mix-Rx-Tx interference, the point-to-point links can also overlap spatially. Because of side-lobe radiation, adjacent directional links need a minimum angular separation to operate. Otherwise, if a node transmits on two adjacent links with insufficient angular separation, then the transmissions will interfere at the receivers [17]. In subsequent discussion we assume that following some topology design process, we are given a network graph that satisfies the angular separation criterion. We return back to the minimum angular separation problem in Section IV.

C. The 2P Protocol

Avoiding Mix-Rx-Tx interference, while keeping a set of links active on the same channel, requires that the subgraph induced by the active links be bi-partite. To see why, consider a subgraph B of the mesh graph that is induced by active links at a given instant. To avoid Mix-Rx-Tx interference in B, every node is either transmitting on all its incident links or receiving on all its incident links. In other words, no two transmitting nodes (and, no two receiving nodes) are neighbors in B. It follows that B is bi-partite.

While Mix-Rx-Tx interference prevents simultaneous transmissions and receptions at a node, a node may synchronously transmit (or synchronously receive) on all its adjacent links. This is called SynTx (or SynRx). SynTx/SynRx together are known as SynOp: synchronous operation of links at a node. In a recent work, Raman et al. [17] have proposed the *2P* (2-Phase) MAC protocol based on SynOp. The protocol operates on bi-partite graphs by switching each node between two phases: SynRx and SynTx. Let $B1$ and $B2$ be the two independent sets of the bipartite graph. To start its transmission, a node v in $B1$ waits for all its neighbors in $B2$ to complete their transmissions. Then, v transmits to all its neighbors in $B2$. (The algorithm for a node in $B2$ is symmetric.) In other words, when a node switches from SynRx to SynTx, its neighbors switch from SynTx to SynRx, and vice versa.

Apart from preventing Mix-Rx-Tx interference, the 2P protocol has two desirable properties: (1) it guarantees 100% utilization of all links (i.e., each link is always active in one direction or the other), and (2) the synchronization is local (i.e., a node only synchronizes with its neighbors). However, 2P has a crucial limitation − it can only operate on a bipartite graph.

D. Routing and Channel Allocation Problem

One approach to schedule links when the given mesh graph is not bipartite is to partition the graph into multiple bi-partite subgraphs. Then, on each bi-partite subgraph, we can run 2P using a different non-interfering channel. However, the number of such channels is limited, e.g., only 3 in 802.11a and g, and we may not always have sufficient channels to cover the mesh graph. Another approach is to activate the bi-partite subgraphs in different time-slots. However, switching from one active bi-partite subgraph to another is difficult without tight time synchronization. We investigate the first approach in this paper.

As we can activate only a bi-partite graph using a single channel, it follows that with K available channels, we can only activate a subgraph that can be covered by K bipartite graphs. However, activating only a subgraph of the mesh graph may degrade the capacity of the network to satisfy node demands. In this case, we need to select K bi-partite subgraphs that best meet the node demands. In Section III, we formally pose this problem and present some initial results.

E. The Topology Construction Problem

One of the primary concerns of VillageNet is to keep the cost of constructing the network as low as possible. Using commodity IEEE 802.11 hardware keeps the cost of communication equipments low. However, to establish a long-distance link, directional antennae need to be mounted on high towers so that a “line-of-sight” can be maintained between the antennae at the two end points of the link. It turns out that the cost of antenna towers is a major component of the cost of the building the network.

The cost of a tower depends on its height. The required height of the tower at a node, in turn, depends on the links of the node. Thus, it is important to judiciously decide which links should be established. To this end, in Section IV, we formulate the topology construction problem: given the locations of the nodes, direct links between which nodes should be established such that, (1) all nodes are connected with a certain robustness in the network, and (2) the cost of constructing towers to establish all the selected links is minimized. We also study the feasibility aspect of this optimization problem: given the locations of the nodes, is it possible to construct towers at the nodes so that robustness requirements of the connected network are met and yet all the constraints in the problem are satisfied?

F. Organization

Before we go forward, here is a quick overview of the organization of this paper. In Section II we discuss some related work in this area. In Section III, we present our approach to select bi-partite subgraphs from a graph G given K non-interfering channels. We also describe a large class of graphs that can be completely covered by K bi-partite subgraphs. We consider node demands, and analyze what fraction of node demands, that can be routed over the entire mesh graph, can also be routed over only the bi-partite subgraphs. Also, through numerical simulations, we evaluate our proposed schemes. In Section IV we describe how to model the antenna heights and costs, the requirements on covering a link, and give a formal description of the topology construction problem. Finally, we summarize and present some concluding remarks in Section V.

II. Related work

VillageNet is motivated by the Digital Gangetic Plains project [2] in and around Kanpur, India. There have recently been other examples of community wireless mesh networks, such as [21] and MIT's Roofnet [7]. However these projects are, for most part, built using omnidirectional antenna, and occupy smaller areas.

Link scheduling in wireless mesh networks is a well-studied problem [10], [9], [6], [20], [4]. However, for our particular setting (long-distance, fixed-topology meshes with directional antennae), the 2P protocol [17] is the only distributed scheduling protocol that we know of. We use 2P as the underlying scheduling protocol in VillageNet.

Our work on channel allocation is close to that of Raman et al. [16]. Raman et al. consider the channel allocation problem in rural mesh networks, and also assume that the 2P protocol will be used for scheduling over bi-partite subgraphs in the given mesh graph. They consider the *specific* case of how to cover the input graph with bi-partite subgraphs when 3 non-interfering channels (as is the case in IEEE 802.11b) are available. They observe that any graph with maximum degree 5 (which by Vizing's theorem is 6-edge colorable) can be covered by 3 bi-partite subgraphs, and propose an algorithm to achieve this based on edge coloring.

We consider the *general* case of how to cover an input graph with bi-partite subgraphs when K channels are available. This can be important since the number of 802.11 channels that are available to the mesh network can vary. For example, 802.11a equipment provides 11 non-interfering channels, as compared to the 3 made available by 802.11b and 802.11g. In some cases, channels may have to be set aside for a local WLAN within a village or because of RF pollution (interference from other neighboring networks). On an opposite note, a recent work [8] has demonstrated the possibility of squeezing out 4 non-interfering channels from IEEE 802.11b and g equipment.

In the case when K bi-partite graphs are not sufficient to cover the entire mesh graph, we also provide guarantees on the fraction of traffic that can be routed over the bi-partite subgraphs, as compared to routing over the entire mesh graph. This is an aspect of our work not covered in [16].

There has also been a considerable amount of work on channel allocation in wireless mesh networks and analyzing how to meet end-to-end demands, for example, in [5], [18], [14], [15]. These works principally differ from ours in that they consider omnidirectional antennae and thus analyze networks with different interference properties as compared to ours.

The work closest to our topology construction problem is [19]. Like us, they also consider the mesh topology construction problem with the antenna tower costs dominating other costs. However, the approach in [19] is to first fix a spanning tree for the network graph, and then formulate the height assignment problem as a LP. We pose the general problem of finding both a spanning tree (or more generally, a spanning subgraph with required connectivity) and height assignment with the lowest costs.

Cellular networks also require efficient schemes to place towers to cover large areas with minimum costs. However the problem in cellular network deployment is to place the minimum number of towers to cover the maximum possible area. In our problem the location of the towers is fixed (namely, the villages), and the goal is instead to select a set of links and height assignments that will cover all nodes with minimum costs.

III. Routing and Channel Allocation

In this section we formulate the routing and channel allocation problem that we discussed in Section I-D. We then give some preliminary results.

Suppose that we are given a mesh graph $G = (V, E)$, where each link is of capacity L and there are K (non-interfering) channels. Then the subgraph that we can activate should be the union of K bipartite subgraphs. (Recall that, given a single channel, the 2P protocol can only activate a bipartite graph.) In addition, suppose that we are given a set of demands $\{(s_i, t_i, d_i)\}$ where (s_i, t_i, d_i) denotes a demand of $d_i > 0$ from s_i to t_i. An immediate problem is to determine what are the best K bipartite subgraphs, corresponding to the K given channels, for meeting these demands.

We pose the above problem in terms of maximum concurrent flow. The *concurrent flow value* of a routing is the largest λ such that the routing satisfies λ fraction of all demands. For a graph G, the *maximum concurrent flow* is defined as the maximum λ such that there is a routing over G with concurrent flow value λ [3].

We would like to activate K bipartite subgraphs such that the max flow of routing on union of these subgraph is as high as possible. If we can cover G using K bipartite subgraph, then we can simply activate the whole graph and compute the max flow for G. (Obviously, for the same set of demands, the max flow for any subgraph of G at most the max flow of G.) Otherwise, we would like to select K bipartite subgraphs such that routing on the union of the K subgraphs provides a guarantee on max flow.

A. Partitioning into Bi-partite subgraphs

To obtain K bi-partite subgraphs of G, we observe that edges in a cut of G naturally specify a bi-partite subgraph. In the following, we will use the term cut and bi-partite subgraph interchangeably. (We will denote the number of edges in a cut C by $|C|$.) We obtain K bi-partite subgraphs by iteratively applying a cut algorithm $\mathcal{C}$. On the original graph $G = (V, E)$, we apply $\mathcal{C}$, and the first bi-partite graph is given by the edges in the first cut. We then remove the edges of the first bi-partite graph from G, and apply the cut algorithm again to obtain the next bi-partite graph. We repeat the process K times to obtain the K bi-partite graphs.

Ideally, we would like our cut algorithm to cover the graph using the smallest number of cuts. As finding a max-cut for a general graph is an NP-hard problem we use the following simple local search based $1/2$ approximation of max-cut [22]. (We denote this algorithm by *LS*.) Initially, the set of nodes are arbitrarily divided into two sets. Each step of the algorithm, called a *flip*, does the following. It selects a node v such that v has more neighbors in its own set than in the other set, and then moves v to the other set. The algorithm terminates when there is no node that can be flipped. It is easy to see that the algorithm takes at most $|E|$ flips to terminate $-$ each flip increases the number of edges in the cut by at least 1. Moreover, upon termination, for each node v, the degree of v in the cut is greater than or equal to its degree in the remaining graph (otherwise, v can be flipped). It follows that on removing K LS-cuts from G, the degree of a node in the remaining graph is at most $1/2^K$ of the original.

We now characterize a class of graphs that can be covered by K LS-cuts, denoted by $B_1, \ldots, B_K$. Let LS_K denote the union these K cuts.

Theorem 1: If the degree of every node in G is at most $2^K - 1$ then LS_K covers G.

Proof: Suppose that the maximum degree of G is at most $2^K - 1$. Consider any node v with degree d. The number of incident edges of v in $G - LS_K$ is at most $d/2^K \leq (2^K - 1)/2^K < 1$. Thus G is covered by LS_K. ■

In particular, in 802.11b and g, where there are three non-interfering channels, Theorem 1 implies that LS-cuts can cover any graph whose maximum degree is at most 7.

B. A Guarantee on Max Flow

We now consider node demands and investigate guarantees on the max flow for a subgraph obtained through K cuts. Recall that, for the same set of demands, the max flow of a subgraph can be at most the max flow of G, and the equality trivially holds when the K cuts cover G (see Theorem 1). Thus we compare the max flow of the subgraph obtained through cuts with that of G.

The following theorem gives a worst-case guarantee on the max flow of a subgraph that is obtained by applying max-cuts, instead of LS-cuts. (See [11] for a proof.) We denote these K max-cuts by $C_1, \ldots, C_K$. Let M_K is the union of these K max-cuts.

Theorem 2: Suppose that at least λ fraction of every demand can be concurrently routed over G. Then $\Omega(\frac{2^K\lambda}{2^K + log(|E|)})$ fraction of every demand can be concurrently routed over M_K.

It immediately follows that if $K = log(log|E|)$ and at least λ fraction of every demand can be concurrently routed over G, then $\Omega(\lambda)$ fraction of each demand can be concurrently routed over M_K.

Theorem 2 gives a worst-case guarantee on the max flow when the bi-partite graphs are obtained using max-cut. Although computing a max-cut is NP-hard in general graphs, it can be computed in polynomial time for planar graphs [13]. Thus, the above theorem gives a guarantee on max flow for planar graphs. Even for general graphs, we now show through numerical simulations on randomly generated graphs that LS-cuts come close to meeting the above guarantee.

We generate graphs for our simulations in a circular plane with a radius of 50Kms. We iteratively place nodes on this plane. For each node we choose the neighbors by randomly selecting nodes within a 10km radius of the node, with a upper bound on the degree of the node. We select one of the node with the greatest number of candidate neighbors to be *gateway* nodes. The capacity of a link is fixed at 11Mbps (the maximum capacity of an 802.11b channel).

For our simulations we create 25 instances of graphs with $N = 75$ nodes. We consider that all nodes have symmetric demands of 8Mbps both to and from the gateway node. We compute the max flow,[1] λ_1, over the *entire* input graph, G. Next, we recursively apply the LS algorithm on the graph G for values of $K = 1, 2, ..., 11$ to identify K bi-partite subgraphs within G. We then compute λ_2, the max flow over the union of these bi-partite subgraphs. We are interested in the ratio $\frac{\lambda_2}{\lambda_1}$. We compute the mean and the standard deviation of this ratio over all graph instances G and plot it in Figure 2 for different values of K. We plot the value of this analytical expression for different values of K. (On average, the number of edges $|E| = 1345$ for these graph instances.)[2] It is clear that the max flow achieved over the union of K-cuts computed using our LS algorithm always meets and exceeds the predicted worst-case performance if we had used max-cuts instead. (In our simulations the average value of λ_1 was 0.33.) Specifically, we observe that for $K = 3$, $\frac{\lambda_2}{\lambda_1}$ is already > 0.95. We also observed a similar trend for variations of this experiment with different node demands, and also for graph instances with two gateway nodes.

IV. The Topology Construction Problem

In this section, we formulate the topology construction problem in VillageNet. Maintaining line-of-sight over long distances require the directional antennas to be mounted on tall towers so as to overcome obstacles such as trees, buildings and terrain. Building these towers can be very expensive. The required height of the towers at either ends of a links depends (among other things) on the length of the link, and the cost of the tower in turn increases in proportion to its height. Table I (taken from [1]) lists the cost of building towers for various heights and materials used.

To establish a link over a distance of 7-8 kms (typical inter-village distance) requires tower heights to be around 30-45 meters [2]. As is evident from Table I, the cost of building such a tower is an order of magnitude greater than the cost of the communication equipment at a node. As mentioned earlier, a primary requirement of building rural networks is to keep the costs as low as possible. Given that the cost of

[1] We formulate the max flow problem as an LP using AMPL [12] and solve it with the CPLEX LP-solver.

[2] Since, we want to evaluate the performance of the LS algorithm (vs. our analysis) for different values of K, we allow the nodes in these graphs to have a maximum degree of 36. This ensures that close to $K = 6$ cuts will be required for the LS algorithm to cover the entire graph.

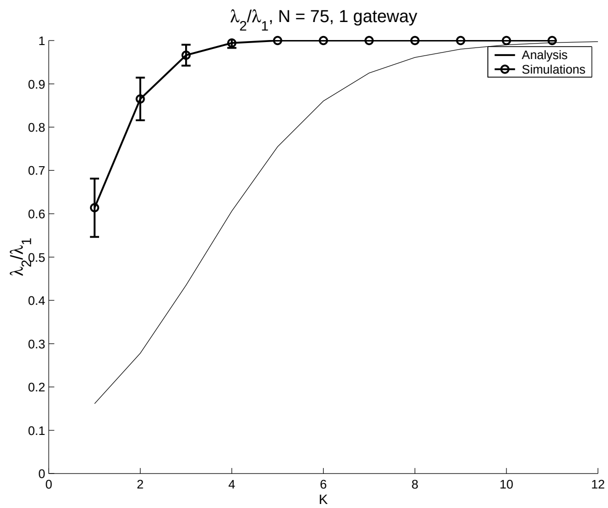

Fig. 2. Ratio of λ_2/λ_1 for different values of K. Comparing the max flow from the LS algorithm with the flow predicted analytically, if max cuts had been used instead.

Items	Cost in US$
Steel tower	
15m	$2,450
20m	$3,100
25m	$3,500
30m	$4,000
40m	$5,100
Antenna mast	
10m	$85
15m	$130
20m	$170
Mesh AP	$150-$200

TABLE I

COST OF ANTENNA TOWERS OF DIFFERENT HEIGHTS, AND OF MESH AP EQUIPMENT AT A NODE

the antenna towers dominates other infrastructure costs, the principal problem in building such networks is to construct a topology with the lowest total cost of building the antenna towers at every node.

Now, let us highlight some of the main issues in this problem.

a) Connectivity Requirements: First, it is important to ensure that the topology constructed must form a connected network. But, what is the exact connectivity that we want in our network? The answer to this question shall depend on several factors. For example, if the network links are known to be unreliable, then we would like to have higher connectivity in the network for the purpose of robustness; on the other hand, if the links are known to be reliable, then it might be advantageous to simply use a spanning tree since, in that case, the network is bipartite and therefore, the 2P protocol can be used. Another criterion that shall influence the connectivity requirements of the network is the value of network flow that we want to achieve. As we have seen earlier, if multiple channels are available for communication, then using 2P protocol does not require the entire network to be bipartite. In that case, we might be interested in achieving higher connectivity (and therefore greater flow) in the network while maintaining degree bounds which shall ensure that the 2P protocol can be used. So, the connectivity required shall depend on multiple factors and has to be provided as an input to the topology construction problem.

b) Physical Limitations: There are multiple physical restrictions on the tower heights as well. For instance, it is prohibitively expensive to build a tower which is higher than some particular threshold value. This height represents the maximum possible height of a tower. Further, antennas placed at a tower need to have a minimum angle of separation between them because their directionality becomes effective only at a significant distance away from the antennas. This effectively means that no two adjacent communicating links should form an angle less than the minimum angle of separation. Another interesting point to note is that the terrain might not be at the same altitude throughout. So, the tower bases could also be vertically displaced with respect to each other.

c) Nature of Cost Function.: An important component in this problem is the nature of the cost function which maps tower heights to the cost of building the tower. As shown in Table I, there are two types of antenna towers. For heights < 15 meters, one can use the cheaper masts. For greater heights, one has to use the more expensive steel towers whose cost increases roughly linearly with height [19]. Further, there is an order of magnitude difference between the cost of the cheaper masts and that of the steel towers. Thus, roughly, the cost function is constant as long as the cheaper masts can be used and becomes linear in height once the steel towers are needed, with a jump in cost when we change from masts to steel towers. This cost function is illustrated in Figure 3. Clearly, this implies that we do not obtain any significant cost savings by varying the height of a tower constructed from the cheaper material. Thus for all practical purposes, we can assume a minimum height of each tower that is equal to the maximum height of an inexpensive masts, which is around 20 meters.

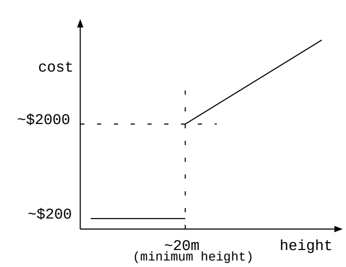

Fig. 3. The cost function

d) Conditions on Tower Heights for establishing a direct link: Finally, we need to establish the conditions on the heights at two towers (mounted at A and B, say) that ensure that a communication link is established between them. Firstly, due to restrictions on the length of a link, nodes that are very distant from each other cannot communicate directly. The proximity of A and B ensures that the transmit powers

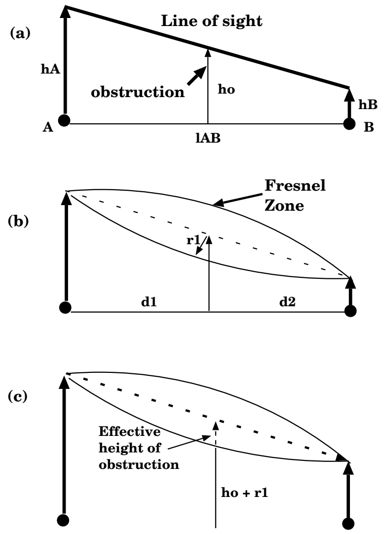

Fig. 4. Computing the height of towers at the end-points of a link

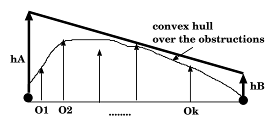

Fig. 5. Computing the height of towers at the end-points of a link with multiple obstructions in between

and the gains of the antennas at both ends are sufficient to overcome the path loss between the two points. Next, we need to establish a line-of-sight between the two antennas if we intend to have a direct link between them. To this end, one needs to ensure that the straight line joining the antennas mounted on the towers should clear any obstructions along the path (as shown in Figure 4a). In fact, it is essential to maintain the stricter *RF line-of-sight* between the two points. This is determined by an elliptical area between the two points, A and B (illustrated in Figure 4b), termed the first *Fresnel zone*. To establish RF line-of-sight, a sufficiently large area of the Fresnel zone should also clear all obstructions between A and B. Consider an obstruction O of height h_o at a distance d_1 from A and d_2 from B. The radius of the first Fresnel at this point is defined as $r_f = \sqrt{\frac{\lambda d_1 d_2}{(d_1+d_2)}}$, where λ is the wavelength of a 802.11b signal. Then, to establish RF line of sight, it is required that at least 60% of r_f be clear of any obstacles at O. In other words, (as shown in Figure 4c) the line joining the antennas at A and B should be at an height $> h_o + 0.6r_f$ at the point O. Another way of looking at this is, that to have RF line-of-sight between A and B, with an obstruction of height h_o at a point O, it would be sufficient to simply establish visual line of sight with an obstruction of height $h_o + 0.6r_f$ at O.

In reality, there maybe be multiple obstructions between A and B. As in Figure 5, consider multiple obstructions, $(O_1, O_2, ..., O_k)$, of heights $(h_{o1}, h_{o2}, ..., h_{ok})$ between A and B. Let the radius of the first Fresnel zone at each of these obstructions be $(r_{f1}, r_{f2}, ..., r_{fk})$, giving the obstructions an *effective* height of $(h_{o1} + 0.6r_{f1}, h_{o2} + 0.6r_{f2}, ..., h_{ok} + 0.6r_{fk})$. Consider a convex hull, defined by a function $O_{AB} : (0, l_{AB}) \rightarrow \mathbb{R}_0^+$ that covers the effective heights of these obstructions. Now, let h_A and h_B represent the tower heights at the nodes of A and B (Assume that both the nodes A and B are at the same altitude; the general case where A and B might be at different altitudes can be easily handled by considering the heights of all towers and obstructions with respect to some common base altitude, as shown later in this section.). Covering edge AB requires a visual and RF line-of-sight connection between the towers at its two terminal nodes. This would imply that the interpolated straight line f_{AB} between the two tower heights should dominate the convex hull of obstructions, O_{AB}. Formally, $f_{AB} : (0, l_{AB}) \rightarrow \mathbb{R}^+$ is defined as

$$f_{AB} = h_A + \left(\frac{h_B - h_A}{l_{AB}}\right) x$$

and $\forall 0 < x < l_{AB}, f_{AB}(x) \geq O_{AB}(x)$.

A. Formal Description of the Topology Construction Problem

Consider the set of villages to be points on a two dimensional plane: if there are n villages, then they are denoted by their Cartesian co-ordinates $\{(x_i, y_i) : 1 \leq i \leq n\}$. Let the following constants be also provided as input:

- $k \in \mathbb{Z}^+$, the required minimum connectivity of the network,
- $\theta \in [0, \pi]$, the minimum angle of separation between any two links incident on a node,
- $h_{max} \in \mathbb{R}^+$, the maximum possible height of a tower,
- $h_{min} \in \mathbb{R}^+$, the minimum height of a tower- recall that this is the maximum possible height of a tower built using the cheaper masts, and
- $L \in \mathbb{R}^+$ is the maximum distance between two points for which the transmit powers and the gains of the antennas at both ends are sufficient to overcome the path loss between the two points.

Let $\mathbb{N}_n = \{1, 2, \ldots, n\}$. Now, let $d : \mathbb{N}_n \times \mathbb{N}_n \rightarrow \mathbb{R}_0^+$ be the Euclidean distance between each pair of points. Thus, $d(i, j) = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}$. Each village also has an associated altitude with respect to some common base level. This altitude is given by the function $a : \mathbb{N}_n \rightarrow \mathbb{Z}_0^+$.

The cost function $c : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is also part of the input. It is defined as

$$c(h) = \begin{cases} K & \text{if } 0 \leq h \leq h_{min} \\ h + K & \text{if } h > h_{min} \end{cases}$$

Here, $K \in \mathbb{R}^+$ is the cost of an antenna constructed using the cheaper masts.

Finally, for each pair of points $(i, j) \in \mathbb{N}_n \times \mathbb{N}_n, i \neq j$, an *obstruction function* o_{ij} is defined. It represents the convex hull of obstructions (taking into account Fresnel clearance) between the points i and j.

Let us consider a height function $h : \mathbb{N}_n \rightarrow \mathbb{R}+$. A link (i, j) is said to be *covered* if the heights of the towers at its two ends (including the altitude at the nodes, i.e. the function $h + a$) are such that a direct communication link can be established between i and j. Let T be the set of links covered by the function $h + a$. Then, our height function must minimize the total cost of constructing the towers at all the nodes, $\sum_{i=1}^{n} c(h(i))$, subject to the following constraints:

- T has an edge connectivity of at least k.
- For any two links in T of the form (i, j) and (i, k), the positive angle between the vectors $\overline{ij}$ and $\overline{ik}$ is at least θ.
- For any point $i \in \mathbb{N}_n$, $h_{min} \leq h(i) \leq h_{max}$.
- For any link (i, j) in T, $d(i, j) \leq L$.
- For any link (i, j) in T, the corresponding obstruction function o_{ij} is cleared by the function interpolated between $h(i) + a(i)$ and $h(j) + a(j)$.

B. Feasibility of a connected Topology

The optimization question described above also poses an associated feasibility question:
Given an input as described above, does there exist a subgraph T such that all the edges satisfy the length and angular separation constraints?
This question is non-trivial. For instance, consider a set of 11 points forming the vertices and center of a regular decagon (refer to Figure 6).

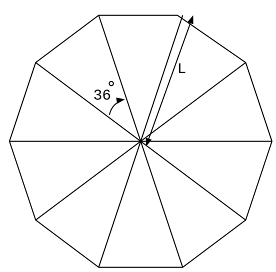

Fig. 6. An example where a feasible solution to the topology construction problem does not exist

Let the maximum link length L be the distance from the center to any vertex of the decagon. Clearly, the only links possible (i.e., of length $\leq L$) are the links between adjacent vertices and links between the vertices and the center. Now, if the connectivity requirement k is 3, the only possibility is to include all the edges in T. However, for a typical minimum angular separation θ of $40°$, including all edges in T violates the angular separation requirement at the central point.

We consider a special sub-case of this problem when the required minimum connectivity $k = 1$. We claim that in this case, any minimum spanning tree (MST) of the given set of points (the distance between points is the weight on the edge joining them) is a subgraph satisfying the angular separation requirements provided $\theta \leq 60°$ (typical values of θ are between $30°$ and $45°$). This follows from the fact that if a MST has an angular separation $< \theta \leq 60°$ at a point i between links (i, j) and (i, k), then the longer of these two links (say (i, j)) can be replaced by the link (j, k) which is shorter, to get a spanning tree with lesser total weight (refer to Figure 7 for an illustration). Also note that $d(j, k) < d(i, j) \leq L$ and therefore (j, k) is a valid link. This contradicts the fact that the original spanning tree was minimal.

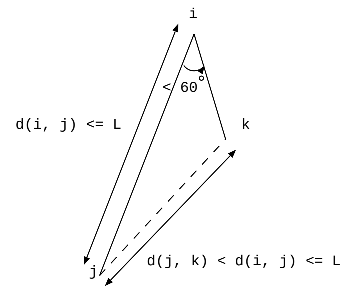

Fig. 7. Two edges in a MST cannot have an angular separation $< 60°$

We are currently investigating the topology construction problem. We have proved a version of the problem to be NP-hard (through reduction from the set-cover problem). We are looking into efficient approximation algorithm, based on a promising greedy technique. This is work in progress.

V. Concluding Remarks

In this paper we gave an overview of VillageNet, a low cost wireless mesh network for rural areas. In VillageNet, nodes are connected using long distance 802.11 wireless links which are established using high-gain directional antennae. The interference patterns in this network imposes the restriction that only a bipartite subgraph of the original mesh graph can be activated using a single 802.11 channel. We studied a natural problem in this setting, given K channels, which K bipartite subgraphs of the mesh graph should be activated to best meet the node demands. We posed the problem as a max flow problem, and showed a guarantee on max flow when the bipartite graphs are obtained using a max-cut algorithm. As max-cut is NP-hard, an immediate question is whether such a guarantee can be provided for bipartite graphs obtained using a polynomial-time algorithm. Another direction of future research is how to activate all links in the mesh graph by scheduling over a set of bipartite subgraphs that cover the whole mesh graph. In this case, switching from one bipartite subgraphs to the other, using the same channel, may involve synchronization issues.

Another direction of research can be how to mask scheduled and unscheduled network failure. Power cuts may be frequent in rural areas. Also, in remote areas, if a link goes down due to equipment malfunction, it might take a long time to replace the faulty equipment. In this setting, it is important to establish a network that remains connected even when few links and nodes go down. Also, if we know in advance when a node will go down, e.g., scheduled power-cuts, then we may want to use this information to improve our routing.

The other main problem we considered is that of determining an optimal set of tower heights so that the necessary connectivity requirement in the network is satisfied while minimizing the cost of topology construction. In this context, we identified all the different characteristics of the problem and modeled the problem mathematically to capture all these features. Finding an optimal solution to the problem remains open. We also studied the corresponding problem of feasibility of constructing a topology satisfying some connectivity requirement and showed that though in general, such a topology might be infeasible, for a connectivity requirement of just 1 and when the angular separation required between adjacent links is at most $60°$, such a topology always exists. This question remains open for the general situation where the connectivity requirement in the network can be arbitrary.

REFERENCES

[1] ANW series heavy duty self supporting towers. http://www.anwireless.com/price.html. 2006.

[2] Digital gangetic plains: 802.11-based low-cost networking for rural areas, 2001-2004: A report. Technical report, IIT Kanpur, 2004.

[3] R. Ahuja, T. Magnanti, and J. Orlin. *Network Flows - Theory, Algorithms and Applications*. Prentice-Hall, 1993.

[4] I. F. Akyildiz, X. Wang, and W. Wang. Wireless mesh networks: a survey. *Computer Networks*, 47:445–487, 2005.

[5] M. Alichery, R. Bhatia, and L. Li. Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks. In *Proceedings of MobiCom*, 2005.

[6] P. Bahl, R. Chandra, and J. Dunagan. SSCH: Slotted seeded channel hopping for capacity improvement in ieee 802.11 ad hoc wireless networks. In *Proceedings of MobiCom*, 2004.

[7] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *Proceedings of MobiCom*, 2005.

[8] M. Burton. Channel overlap calculations for 802.11b networks. 2002. http://www.80211-news.com/publications/page289-376.asp.

[9] I. Chlamtac and S. Kutten. A spatial reuse tdma/fdma for mobile multi-hop radio networks. In *Proceedings of Infocom*, 1985.

[10] I. Chlamtac and S. Lerner. A link allocation protocol for mobile multi-hop radio networks. In *Proceedings of Globecom*, 1985.

[11] P. Dutta, S. Jaiswal, and R. Rastogi. Routing and channel allocation in rural wireless mesh network. In *Proceedings of Infocom, to appear*, 2007.

[12] R. Fourer, D. M. Gay, and B. Kernighan. *AMPL A modeling language for mathematical programming*. Thomson, 2 edition, 2003.

[13] F. O. Hadlock. Finding a maximum cut of a planar graph in polynomial time. *SIAM Journal on Computing*, 4:221–225, 1975.

[14] M. Kodialam and T. Nandagopal. Characterizing achievable rates in multi-hop wireless mesh networks: the joint routing and scheduling problem. In *Proceedings of MobiCom*, 2003.

[15] P. Kyasanur and N. Vaidya. Capacity of multi-channel wireless mesh networks. In *Proceedings of MobiCom*, 2005.

[16] B. Raman. Channel allocation in 802.11-based mesh networks. In *Proceedings of IEEE Infocom*, Apr. 2006.

[17] B. Raman and K. Chebrolu. Design and evaluation of a new MAC for long distance 802.11 mesh networks. In *Proceedings of ACM Mobicom*, Aug. 2005.

[18] A. Raniwala and T.-C. Chiueh. Architecture and design for an IEEE 802.11-based multi channel wireless mesh network. In *Proceedings of Infocom*, 2005.

[19] S. Sen. Topology planning for long distance wireless mesh networks. Master's thesis, IIT Kanpur, 2006. http://www.cse.iitk.ac.in/users/braman/students/2006/sdsen.html.

[20] J. So and N. Vaidya. Multi-channel MAC for ad hoc networks: handling multi-channel hidden terminals using a single tranceiver. In *Proceedings of Mobihoc*, 2004.

[21] R. van Drunnen, J. Koolhaas, H. Schuurmans, and M. Vijn. Building a wireless community network in the netherlands. In *Proceedings of Usenix/Freenix conference*, 2003.

[22] V. Vazirani. *Approximation Algorithms*. Addison-Wesley, 2001.