Sparse Multi-hop Wireless for Voice Communication in Rural India

Srinath Perur and Sridhar Iyer
KReSIT, IIT Bombay, India
e-mail: [srinath,sri]@it.iitb.ac.in

Abstract

There are many villages in India with sizeable populations that have no means of using telephony for communication within the village. We identify *sparse* multi-hop wireless networks carrying packetised voice as a quick, cost-effective and flexible solution for rural telecommunication. We discuss design considerations for such a network and define a metric called reachability that is useful in evaluating design tradeoffs in sparse networks. We present a case study where we use simulation to: i) answer specific questions related to deployment; and ii) obtain insights into the nature of design tradeoffs in such networks.

1 Introduction

The Department of Telecommunications (DoT), through its Village Public Telephone (VPT) scheme, aims to have at least one telephone installed in each of approximately six lakh villages identified in the 2001 census [1]. As of August 2005, VPTs have been deployed in 83.3% of the targeted villages [2]. The next phase involves installing a second telephone in villages with a population over 2000. The current focus of rural telecom initiatives is rightly to connect villages to the world outside. At the same time, there is also a need to connect people *within* a village. Census figures show that around half of all Indian villages have populations between 500 and 2000. Since these villages are predominantly agricultural, their inhabitants are spread over fairly large areas making local communication desirable. But neither cellular nor fixed telephony is likely to be viable in several villages for some time to come. This is due to the service providers' inability to recover infrastructure costs, and is borne out by statistics which show that cellular coverage in Indian rural areas is negligible at present [3]. There are several efforts being made to bring connectivity to villages. Besides DoT and TRAI (Telecom Regulatory Authority of India) schemes, WLL (Wireless in Local Loop) solutions using corDECT [4] and the Digital Gangetic Plain project [5] are recent initiatives to connect villages to the world outside. In addition to these, we believe efforts are required to find ingenious ways to connect people *within* a village.

A possible means for enabling local communication within rural areas is through deploying multi-hop wireless networks that carry packetised voice. Individuals would carry inexpensive hand-held devices capable of encoding/decoding voice and performing multi-hop routing. These devices would form a network that facilitates communication in two modes: i) real-time VoIP conversations; and ii) offline voice messages. The offline voice messaging mode would be used when the network cannot satisfy bandwidth and connectivity requirements for a real-time conversation, and it can be used to communicate asynchronously using store and forward mechanisms. Such a system has several advantages in the rural scenario: it does not require any infrastructure deployment apart from the hand-held devices themselves, and as a result is relatively inexpensive and quick to deploy. This also makes it possible to use these networks as a short term arrangement while other efforts are underway to bring connectivity to villages. Such a system also does not have a single point of failure, is robust, and degrades gracefully. This is an advantage where regular system maintenance cannot be guaranteed. Enabling communication in remote areas is a well known application for wireless ad hoc networks, but deploying *sparse* networks in constrained application scenarios is not very well studied. Such an approach introduces an additional degree of flexibility: we can trade deployment cost for performance depending on the application's requirements and the available resources. Understanding how to evaluate this tradeoff is critical to having useful deployments of sparse multi-hop wireless networks.

In this paper we identify sparse multi-hop wireless networks as being a possible means for local communication in the Indian rural context (Sec. 2). We define a connectivity property called *reachability* which is a sensitive measure of the extent of communication supported by a sparse network (Sec. 2.2). We use reachability to explore a case-study through simulations, and evaluate tradeoffs between parameters such as transmission range, number of devices, area of deployment, and communication capabilities of the network (Sec. 3).

2 Design Considerations

In designing a multi-hop wireless network, some of the following parameters may be known or given, and some will have to be decided upon by the designer: the number of devices, capabilities and cost of each device, dimensions and topography of the deployment area, usage pattern, and level of connectivity desired in the network. If the deployment is a dense one, interference between nodes, and the resulting loss in network capacity must also be considered. These design parameters are also interdependent to a large degree. Increasing the number of nodes is likely to increase connectivity, but this also increases the cost of deployment. Trying to ensure the same level of connectivity while using fewer nodes would require us to increase transmission range. A small increase in transmission range could result in a large increase in the power consumption of a node. This would result in either a shorter life for nodes, or a need for more expensive nodes with batteries of higher capacity. Transmission range also depends on the physical terrain in the area of deployment: the same transmission power would result in a longer range in a flat, field like area, and a shorter, fluctuating range in the presence of uneven, wooded terrain. Multi-hop ad hoc networks are also known to exhibit *phase transition* behaviour—a small change in transmission range or the number of nodes can cause large changes in connectivity properties [6]. When nodes are capable of movement, the speed and pattern of mobility, and their effect on network performance must also be considered.

2.1 Sparse networks

Connectivity is defined as *the probability that all nodes in the network graph form a single connected component*. The factors affecting connectivity are primarily the dimensions of deployment, the number of nodes within it, and the nodes' effective transmission range. If these three parameters characterised a network, some of this network's instances would be connected, while others would not. The connectivity of this network can be measured as the fraction of network instances in which the network graph is completely connected. For a network with mobile nodes, connectivity can be expressed as the fraction of the network's lifetime during which it was completely connected.

A *sparse* network is one that is unlikely to be fully connected. An important design consideration is deciding how much connectivity is acceptable. Complete connectivity may always be desirable, but may not be achievable at an acceptable cost. In such cases we may be willing to tolerate a lower level of connectivity. There is work that shows that an ad hoc network willing to tolerate a small degree of sparseness can use a transmission range much lesser than that required for full connectivity [7]. Similarly, a sparse network would also need substantially fewer nodes for slightly reduced connectivity. Using fewer nodes or a smaller transmission range translates into lower deployment costs. This ability of sparse multi-hop networks to trade cost for connectivity makes them particularly well-suited for economically constrained rural deployments.

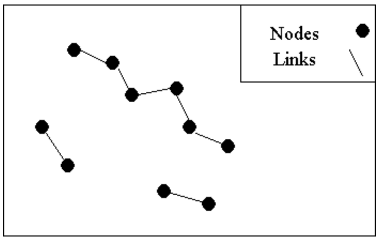

Figure 1: A network instance with reachability=0.378

2.2 Reachability

In densely connected networks, it is almost certain that one node can reach another. In fact much work in the area of ad hoc and sensor networks focuses on determining the conditions that will ensure a completely connected network. In our case, we would like to deliberately design a sparse network that trades connectivity for cost. To accomplish this, we need a fine-grained indicator of the extent to which a network is connected.

The reachability of a static network is defined as the *fraction of connected node pairs* in the network. It is a property of the network graph, with no assumptions made regarding the distribution of nodes. Using this definition we can calculate reachability for a network of N nodes as:

$$Reachability = \frac{\text{No. of connected node pairs}}{\binom{N}{2}} \tag{1}$$

Like connectivity, this too depends on the dimensions of deployment, the number of nodes, and the nodes' effective transmission range. A pair of nodes is considered connected if there is a path of length one or greater between them. Figure 1 shows one instance of a network with 10 nodes. We count the number of node pairs that can reach each other, that is, nodes that are connected either directly or through other nodes, as 17. Substituting $N = 10$ in the denominator of Eqn. 1, we obtain the reachability for this network instance as 17/45 or 0.378.

Note that for the same 10 nodes, it is possible to have a different value of reachability in another instance. A network for our purposes can be defined by the number of nodes, their bounding area, and the transmission ranges of the nodes. When we speak in general of the reachability of a network being c, we imply *probabilistic* reachability—if many network instances are observed, and their reachabilities measured, the long term mean would tend to c. This expected value is significant since it represents the probability that a node pair chosen at random from the the network is connected. Similarly, note that a single instance of a network is either fully connected or not, and connectivity is measured as the fraction of a large number of network instances that are connected.

When nodes are mobile, the fraction of connected node pairs varies from time to time depending on node movement, but a single value can be obtained for any time *in-*

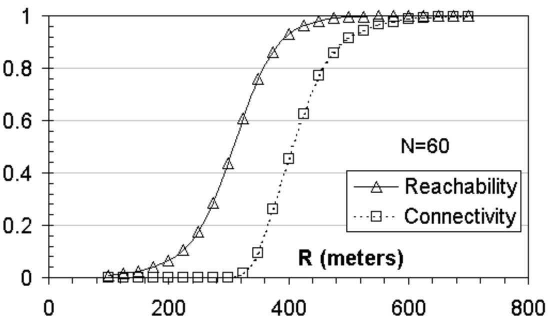

Figure 2: Reachability and Connectivity vs. R

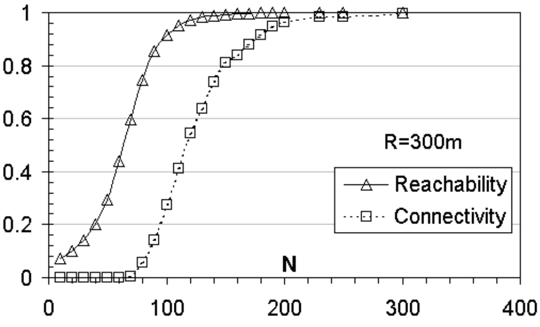

Figure 3: Reachability and Connectivity vs. N

stant. We can measure reachability for a mobile network as the *average of instantaneous reachability values* measured at frequent intervals during the operation of the network.

While studying the rural deployment scenario, we prefer using reachability to connectivity since it is: i) a more intuitive measure of the extent of communication possible between pairs of nodes; and ii) more sensitive to changes in the number of nodes or transmission range, especially for sparse networks. More details about reachability, including a regression model to help designers evaluate tradeoffs can be found in [8].

3 Case Study

Consider a village with a few hundred inhabitants that is spread over an area of 2km x 2km. Quite a large portion of the village is agricultural land, contributing to the low density of inhabitants. A number of devices capable of multi-hop packetised voice communication are to be deployed among people in the village. We now identify some design tradeoffs in this scenario through extensive simulations.

3.1 Simulation Preliminaries

The simulations presented in this paper are conducted using Simran [9], a simulator we have developed for studying topological properties of multi-hop wireless networks. Simran takes as input a scenario file with initial positions and movements of nodes, and generates a trace file containing metrics of interest such as average number of neighbours, averaged shortest path lengths over all pairs of nodes, reachability, connectivity, and number and size of connected components. Simran is also supported by a number of smaller programs for generating scenario files, managing simulations and for analysing results. Simran also supports topological simulation of networks with asynchronous communication.

Initially, in sections 3.2, 3.3, and 3.4, we assume that mobility is low enough that when a connection exists between two nodes, it is unlikely to break while a call is in progress. We treat the network as static with N nodes distributed uniformly at random over the area of operation. Later, in sections 3.5 and 3.6, we relax this assumption for assessing the impact of mobility and asynchronous communication. Usually transmission range depends on the transmission power at each node, terrain, presence of structures that cause radio interference, and antenna characteristics at the receiver. For simplicity, we assume that all nodes have a uniform transmission range, R.

3.2 Choosing R

- If there are 60 devices available for deployment in the village, and each device has a transmitter with power control, from what range of values should R be chosen?

To answer this question, a graph such as Fig. 2 is useful. It plots reachability and connectivity against R for 60 nodes. Each point on the graph is the average of 500 simulations. The graph tells us, for instance, that setting the value of R at 100m will certainly not facilitate communication in the village. Similarly, setting R to a value above 600m is unnecessary since the network is already fully connected at that point. We can set the value of R for any desired value of reachability or connectivity. However, as R is set higher, there will be a corresponding increase in the node's power usage. When R is in the phase where the network's connectivity or reachability is growing rapidly, small changes in R can result in large changes in the extent to which the network is connected.

3.3 Choosing N

- R is fixed at 300m for a specific device's capabilities in the local terrain. How many nodes are required to be operational in order to ensure that a villager who tries to make a call to another succeeds on average 60% of the time?

This question can be answered from finding the value of N corresponding to a reachability of 0.6. From the graph in Fig. 3, we learn that we would need around 70 devices operational in the area. (Note also that Fig. 3 provides

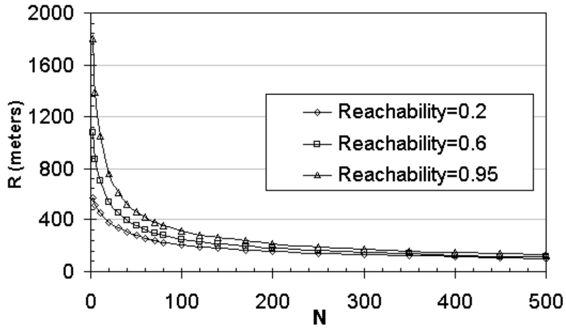

Figure 4: R vs. N for different values of Reachability

an illustration of our claim in Sec. 2.2 that reachability is more sensitive than connectivity for sparse networks. When reachability is 0.6, the corresponding value of connectivity is not useful since it is still at zero.)

An interesting observation can be made from Fig. 2 regarding the relationship between the two measures used—reachability and connectivity. With R set to 400m, reachability is almost at 1, but connectivity does not reach 1 till R is around 600m. This implies that the extra 200m required to ensure full connectivity contributes very little towards increasing the number of node pairs that can communicate. At the same time, the extra range comes at a very high cost since transmission power varies as a power law function of distance.

3.4 R vs. N

Figure 4 shows the relationship between the values of R and N required to keep reachability fixed at 0.2, 0.6 and 0.95. Note that as N decreases below a threshold, the required value of R increases steeply. Given the maximum value R can take for a device, we can find the minimum number of those devices required to be operational for achieving the required reachability. As the network evolves, more nodes may join, or some nodes may be switched off. If we are in a position to implement distributed power control at the nodes, we can use curves like these to maintain reachability at a desired level.

In [8], we characterized reachability using an empirical regression model. We have used this model to build a design tool for sparse multi-hop wireless networks: given three values from deployment area, reachability, R and N, it calculates the fourth. Figure 4 has been generated using this tool.[1]

3.5 Coverage

Since the network we are studying is sparse, we are interested in investigating if nodes are connected only to nearby nodes. If all the node pairs that contribute to the reachability of the network are located near each other, then the network would only be facilitating communication between people who are already within easy reach. To see the extent of a path's coverage, we find the following theorem useful.

Theorem 3.1. *Let $G = (V, E)$ be a graph in which every pair of nodes $(u, v) \in V \times V$ is assigned a distance $|uv|$, and $(u, v) \in E$ iff $|uv| \leq R$. Then, if the shortest path between some two nodes in V has k edges, $k > 1$, the sum of the distances of those k edges, L, is bounded as: $\lfloor \frac{k}{2} \rfloor R < L \leq kR$.*

Proof. The upper bound is trivially kR. $L > kR$ would imply at least one of the k edges being larger than R, which is not possible by definition.

When $k = 2$, let nodes u, v, w in order be the nodes on the shortest path. Then, $L = |uv| + |vw|$ cannot be less than or equal to R since this would imply $(u, w) \in E$. This is clearly not possible since the nodes u, v, w define a shortest path. Therefore $L > R$ when $k = 2$. When $k = 4$ with a shortest path defined by nodes u, v, w, x, y in order, $|uw| > R$ and $|wy| > R$, implying $L > 2R$. Similarly extending this argument for all even k, $L > \frac{k}{2}R$. This same lower bound must also hold for the shortest path of odd length $k + 1$, since adding an edge cannot *decrease* L. Therefore, for all $k > 1$, $L > \lfloor \frac{k}{2} \rfloor R$. □

We ran simulations with $N = 70$ and $R = 300$, and averaged the length of the shortest path between every pair of connected nodes. The maximum value we saw in any of the 500 simulated network instances was 9.24, the minimum was 2.01, and the average shortest path length was 5.24. From the above theorem, an average shortest path length of around 5 implies a piece-wise linear distance greater than 600m, and at most 1500m in the average case. This indicates that the network is capable of connecting pairs of nodes that are not necessarily located near each other. The mean reachability observed was 0.6.

3.6 Mobility

To investigate the effect of mobility, with $N = 70$ and $R = 300$, nodes were made to move at a speed between 0.5ms^{-1} and 2ms^{-1} following the random waypoint mobility model. The simulation time was 12 hours in which nodes moved to random destinations, paused for half an hour, and then continued moving to another random destination. This mobility pattern was chosen to approximate the movement of people across one day. We found that reachability had increased to 0.71 from the value of 0.6 observed for the static network. This is consistent with our experience that mobility typically improves reachability in a sparse network.

[1]Available from http://www.it.iitb.ac.in/~srinath/tool/rch.html

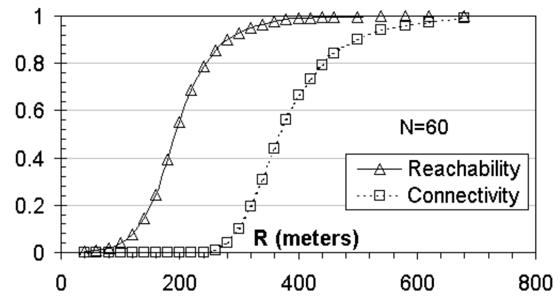

Figure 5: With asynchronous communication

3.7 Asynchronous Communication

Asynchronous communication is particularly useful in sparse networks when routes are difficult to find between source and destination. A message may be passed on to other nodes in the vicinity of the source, and these nodes in turn propagate the message till it reaches the destination. Thus, a message may travel from source to destination without a complete path existing between them at any time. Message Ferrying [10] and routing in delay tolerant networks [11] are representative examples of such asynchronous communication.

We extended the scenario from Fig. 2 to include some degree of asynchronous communication. R was varied keeping $N = 60$. Nodes moved at a uniform velocity of $5ms^{-1}$ without pause. For purposes of calculating reachability, a node pair was considered connected at simulation time t if a path, possibly asynchronous, existed between the two nodes within $t + 30$ seconds. This translates to asking whether a packet with a timeout of 30 seconds can be successfully transmitted between the two nodes using a store and forward mechanism. Similarly, for connectivity, the network was considered connected at a time instant t if all nodes could reach each other asynchronously within time $t + 30$. Averaged values of 20 simulations of 500 seconds each are shown in Fig. 5. On average, nearly 80% of node pairs are connected before connectivity begins to increase from zero. This indicates that sparse networks can achieve a significant degree of communication by operating asynchronously.

4 Conclusions

In this paper, we identified sparse wireless multi-hop networks as being a possible means for facilitating telecommunication within villages in India. We discussed design considerations for such networks, and defined a metric called reachability to help evaluate design tradeoffs. We presented a case study and analysed it through extensive simulation. The broad conclusion of this study is that we can achieve a substantial degree of communication by deploying sparse multi-hop wireless networks. The extent of communication achieved is even more significant when the sparse network is capable of mobility or asynchronous communication. While this case study made simplifying assumptions, we believe it provides interesting insights into the *nature* of tradeoffs involved in designing wireless multi-hop networks for use in rural areas.

References

[1] Department of Telecommunications, India. Guidelines for Universal Service Support. *Downloadable from http://www.dotindia.com/uso/usoindex.htm*, 2005.

[2] Department of Telecommunications, India. Village Public Telephone—Database and Monitoring System. *http://www.moc.gov.in/report.asp*, 2005.

[3] Consultation paper by TRAI. Growth of telecom services in rural india. *Downloadable from http://www.trai.gov.in/consultation.htm*, October 2004.

[4] corDECT. Wireless Access System. *Technical Report, Midas Communication Technologies, Pvt. Ltd.*, 2000.

[5] Pravin Bhagwat, Bhaskaran Raman, and Dheeraj Sanghi. Turning 802.11 inside-out. *ACM SIGCOMM Computer Communication Review*, 34(1):33–38, 2004.

[6] Bhaskar Krishnamachari, Stephen B. Wicker, and Ramon Bejar. Phase transition phenomena in wireless ad hoc networks. In *Symposium on Ad-hoc Wireless Networks, IEEE Globecom*, 2001.

[7] Paolo Santi and Douglas M. Blough. The critical transmitting range for connectivity in sparse wireless ad hoc networks. *IEEE Transactions on Mobile Computing*, 2(1):25–39, 2003.

[8] Srinath Perur and Sridhar Iyer. Characterization of a connectivity measure for sparse wireless multihop networks. *Technical Report, downloadable from http://www.it.iitb.ac.in/∼srinath/rch.pdf*, 2005.

[9] Simran. Simulator for studying topological properties of multi-hop wireless networks. *http://www.it.iitb.ac.in/∼srinath/simran/*.

[10] W. Zhao, M. Ammar, and E. Zegura. A message ferrying approach for data delivery in sparse mobile ad hoc networks. In *MobiHoc '04: Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing*, pages 187–198, 2004.

[11] Sushant Jain, Kevin Fall, and Rabin Patra. Routing in a delay tolerant network. In *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 145–158, 2004.