

FRACTEL: A Fresh Perspective on (Rural) Mesh Networks

Kameswari Chebrolu and Bhaskaran Raman
Indian Institute of Technology, Kanpur
{chebrolu, braman} AT iitk dot ac dot in

ABSTRACT

The use of commodity 802.11 hardware to provide network connectivity to rural regions is an appealing proposition. In this paper, we consider such networks, with a combination of long-distance and short-distance links. In such a setting, we offer a fresh perspective on a variety of technical issues in multi-hop mesh networks. To support QoS for voice, video-based real-time applications, the use of a TDMA-based MAC is appropriate. In this context, we argue that existing approaches to TDMA scheduling and channel allocation are either inapplicable, or are too general and hence complicated. We apply extensive domain knowledge in designing a solution applicable in our context. We also suggest appropriate implementation strategies for the TDMA MAC, capable of scaling to large networks. In all of the above topics, we articulate open technical issues wherever applicable.

1. INTRODUCTION

Mesh networks based on IEEE 802.11 [1] are a cost-effective option for providing Internet connectivity to rural regions, especially in developing countries [2, 3]. In FRACTEL¹, we consider networks having a combination of (a) *long-distance* links: a few km to few tens of km, as well as (b) *local-access* links: shorter distance links, up to about 500 m. The long-distance links extend connectivity from a point of wired connectivity, called the *landline* to a specific point in each village [2]. The local-access links extend connectivity from this point, which we term the *local-gateway*, to multiple nearby locations². Such nearby locations may include individual buildings such as a school, health centre, community centre, residential houses, etc.; it may also include nearby villages. An example deployment setting, with reference to the Ashwini network [4], is shown in Fig. 1.

In such a setting, we wish to support a variety of applications: http/ftp, as well as voice, video-based real-time applications. The Ashwini project has already demonstrated a need for such support: applications like remote education, tele-medicine, agricul-

¹wiFi-based Rural data ACcess and TELephony

²The network thus has a fractal-like structure!

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NSDR'07, August 27, 2007, Kyoto, Japan.

Copyright 2007 ACM 978-1-59593-787-2/07/0008 ...\$5.00.

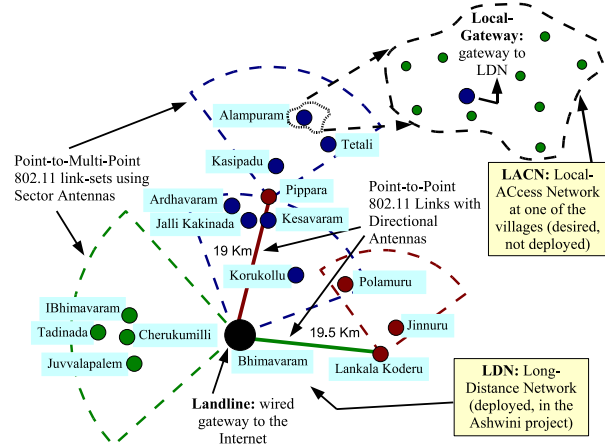


Figure 1: Long-Distance Network (LDN) and Local-Access Networks (LACNs) in Ashwini

tural lessons, etc. based on video-conferencing are the cornerstone of Ashwini [4]. To support such applications, we need to provide adequate performance, with the Quality-of-Service (QoS) necessary for voice/video. We seek to address the challenge of providing this for a large network, consisting of say, a few hundred nodes within a district³.

The basis for enabling QoS is to first have individual links with predictable performance. In prior work [5], we have shown that the *link abstraction* can be made to hold in long-distance links, in the absence of interference. By this, we mean that we can build links to operate at a given data-rate (e.g. 11 Mbps in 802.11b), and with close to 0% error rate, based on an RSSI threshold.

More recently, we have shown the same result to hold for FRACTEL's local-access links too⁴. This is in sharp contrast with what is known from measurements in urban mesh networks; the study of Roofnet in [6] reports that most links have error rates neither close to 0% nor close to 100% (i.e. with intermediate loss rates). It hence concludes that the link abstraction does not hold. The reason attributed in [6] for the loss of the link abstraction is multi-path induced delay spread. Using our own measurements, and also using a fresh analysis of the data from [6], we show that external *interference*, and not multi-path induced delay spread, is the cause of error rates in such links⁴. And in contrast to urban settings, we expect such interference to be uncommon in rural settings.

³Most districts in India are within 20-30 km in radial length.

⁴These measurements are reported in detail in a writeup which is currently under submission.

To achieve QoS in the network, we need to build on the above link abstraction, and have an appropriate MAC scheme in the network. For QoS guarantees, a TDMA-based MAC approach is appropriate. This gives rise to the problem of time-slot allocation (scheduling) as well as channel allocation for the links. Our first contribution in this paper is the articulation of novel approaches to solving these problems (detailed in Sec. 3). While literature is abound with solutions to these issues in mesh networks (see [7, 8, 9] and references therein), we argue that the problem is unduly complicated when considered in a *generic* setting. We apply extensive domain knowledge to identify constraints unique to our setting. We show that channel and time-slot allocation can be considered independently in the Long-Distance Network (LDN), and in each of individual Local-Access Networks (LACNs). And we also suggest approaches for how each of these may be solved.

Our next contribution is with respect to the implementation of the TDMA scheme (detailed in Sec. 4). We are faced with the significant challenge of scalability. We propose four related implementation strategies toward this: (a) the use of hierarchy in the network, (b) centralized scheduling or allocation, (c) a multi-hop, connection-oriented link layer, and (d) a multi-hop framing structure in the LACNs. These strategies fit in well with one another in the envisioned FRACTEL deployment setting. We explain how these strategies address the implementation challenges, and identify open issues wherever applicable.

Existing literature on 802.11 mesh networks for rural use has mainly focused on the use of only long-distance, point-to-point links [10, 5, 3]. And even in these settings, achieving scalable QoS, with considerations of time-slot scheduling and channel allocation is an open issue. FRACTEL considers a more generic setting, with long-distance (point-to-point and point-to-multipoint) as well as local-access links. And we address the issue of achieving QoS scalably.

2. FRACTEL: PROBLEM SETTING

We now discuss two aspects of the envisioned FRACTEL deployment setting, relevant to the later sections: the network architecture, and the nature of the traffic.

Network Architecture: As mentioned earlier, we consider a network which consists of *long-distance* links as well as *local-access* links. FRACTEL makes a specific architectural distinction between the two types of links: while long-distance links involve significant infrastructure and planning, local-access links do not. This distinction surfaces in two main aspects of the network architecture: (a) the types of antennas used, and more importantly, (b) the heights at which antennas are mounted. We explain these with reference to the example in Fig. 1.

Antenna type: Long-distance links can be a few km to few tens of km in length. Such links are formed with the use of high-gain directional antennas. It is common to use antennas with gains of 24-27 dBi and beam-widths of about 8° . Such high gain is necessary to achieve the long range. Links formed with such antennas are typically *point-to-point* (P2P) links. It is also common to use a sector antenna at one of the ends, to construct a *point-to-multipoint* (P2MP) link-set. Examples of both P2P links and P2MP link-sets are shown in Fig. 1.

A sector antenna used in a P2MP link-set typically has a gain of about 17-19 dBi, and a beam-width of $30^\circ - 90^\circ$. A P2MP link-set consists of several links, each between a central sector antenna, and a directional antenna at the other end. The central sector antenna merely replaces the functionality of several directional antennas and circumvents the need for mounting an inordinate number

of antennas at the central location. This of course comes at the cost of significant reduction in antenna gain.

In contrast to such long-distance links, local-access links typically use omni-directional antennas (8 dBi gain). Or they may use low-cost *cantennas*, with a gain of up to 10 dBi. As opposed to high-gain directional or sector antennas, these omni antennas or cantennas are much lower in cost (U.S.\$10-15 or lower, as opposed to \$100 or more). They are also lighter weight and hence easier to mount. Furthermore they avoid the alignment procedures required in using directional antennas. For these reasons, they are more appropriate for individual buildings/houses in the LACN.

Antenna mounting height: A more significant difference between the two types of links is the following. Long-distance links are usually formed by using high-rise towers: 25-50 m tall. Such towers are a significant part of the infrastructure and constitute the major portion of system cost [4]. The high-rise towers are required for line-of-sight and Fresnel clearance above trees and other obstructions. This is depicted in Fig. 2.




Figure 2: Antenna mounting: long-dist,local-access links

It is very advantageous to have a tall central tower (45-50 m), and have it form links with several other shorter towers (25-30 m). This reduces the overall system cost [11], which is a primary consideration in rural networks. In some instances, it is possible to form a long-distance link (10-15 km) between a tall tower (45-50 m) on one end and a *mast* on the other end [11]. A mast is nothing but a metal pipe, up to about 15 m tall.

In contrast to this, antennas for local-access links are mounted on buildings, trees, or other such structures. The mounting height will likely be 5-10 m at most, especially in rural locations with few high-rise buildings. This avoids the cost of tall towers in the LACN.

Network expanse: Apart from the above distinction between the types of links, we wish to draw attention to the expected expanse of the network. Typically, we have one point of wired connectivity in each district, with a radius of about 20-30 km [12]. This suggests that we should be able to cover most districts within two hops of long-distance links from the landline, although we do not rule out instances where more than two hops are necessary.

Nature of Traffic: The next important consideration is the expected traffic. As mentioned, we wish to support voice, video-conferencing, and http/ftp applications. We expect two kinds of traffic to dominate in the network. (1) Traffic between the landline node and another node in the network: video-conferencing sessions between the central location and village nodes are the most common mode of network usage in the Ashwini project [4]. Such traffic is generated within the FRACTEL network, and does not involve the rest of the Internet. (2) We also expect traffic between village nodes and the rest of the Internet, via the landline node.

Specifically, we expect, at least as of now, traffic between two village nodes in the network to constitute only a small fraction. We

expect the bulk of the traffic to be from/to the landline node. We hence optimize for this case.

3. TDMA IN FRACTEL

As mentioned earlier, we have shown in parallel work that we can achieve the *link abstraction*; to have links operate at close to 0% error rate, both in the LDN as well as in LACNs. Building on this predictable performance achieved for individual links, we next need an appropriate MAC capable of providing delay guarantees in the network. This is the topic of discussion in the rest of the paper.

It is well known that multi-hop CSMA/CA performs poorly in terms of achieved capacity as well as the lack of delay bounds. There is a wide body of literature which considers a TDMA-based MAC for multi-hop mesh networks (see [7, 8, 9] and references therein). In literature, typically two related issues are considered: time-slot allocation (TDMA scheduling), and channel allocation for links.

For each link, we need to allocate a (ts_i, c_j) tuple, where ts_i is a time-slot in the TDMA schedule, and c_j is a channel of operation. We assume that we have a set of non-overlapping channels to work with (802.11b/g has three such channels). For any two links which are mutually interfering, we cannot allocate the same (ts_i, c_j) tuple. An optimal schedule minimizes the number of time-slots used.

In such formulation, the problem of (ts_i, c_j) allocation for links in the network translates to the node colouring problem in the corresponding *interference graph*. And optimal node colouring in general graphs is a well known NP-hard problem. The focus in literature has thus been on developing approximate algorithms and heuristics in these settings. Recent formulations consider routing in the mesh also to be a variable [7, 8]. The expected traffic pattern is also taken as an input to the problem. And so is the number of radios available at each mesh node. With such additional aspects included, the problem becomes further difficult.

We argue that in our setting a general consideration is not warranted. We now show how we can apply domain knowledge to identify unique constraints specific to our setting, and simplify the problem considerably.

Spatial reuse in FRACTEL

A generic consideration is inappropriate in FRACTEL primarily because the network has a lot of (hierarchical) structure to it. More importantly, the spatial reuse possible also has a definite structure. We start with the following subtle insight. **Observation O1:** *the LDN, and the LACNs at each village, are independent of one another.* That is, they are non-interfering.

To see this, recall from Sec. 2 that a LACN consists of antennas mounted at short heights (5-10m height). Also recall that a LACN consists of relatively lower gain antennas (8-10 dBi). These mean that the local-access links are really “local”, and the LACNs of different villages are going to be non-interfering with one another.

And each LACN would also be non-interfering with the LDN, but for one exception: each LACN has a local-gateway node connecting to the LDN. The long-distance links at the local-gateway may interfere with the LACN. For instance, in Fig. 1, the long-distance link Pippara-Alampuram may interfere with the local-access links in Alampuram’s LACN. This exception too can also be resolved easily, as we show below.

O1 above considerably simplifies the problem, and reduces its scale: we can now consider time-slot, channel allocation independently in the LDN, and independently in each LACN. We now address these.

Time-slot, channel allocation in the LDN

It is known that the problem of routing is related to that of (ts_i, c_j) allocation in a generic mesh network [7]. We discuss this aspect first. We first note that given that most traffic is from/to the landline (Sec. 2), it is natural to consider a routing tree rooted at the landline. More importantly, multipath routing opportunities are few in the LDN. This is because we typically use directional antennas for better gain, to connect toward the landline node (e.g. in Fig. 1, Juvvalapalam uses a directional antenna toward Bhimavaram). This means that the topology itself has a natural tree structure. This then suggests our next observation **O2:** *since the LDN topology is naturally a tree, the issue of routing can be ignored during time-slot, channel allocation.*

Proceeding further, recall from Sec. 2 that many practical scenarios involve just two hops from the landline. So let us first consider trees of depth two. Call the nodes one hop from the landline as hop-1 nodes, and similarly nodes two hops from it as hop-2 nodes. Hop-1 links connect hop-1 nodes to the landline, and hop-2 links connect hop-2 nodes to hop-1 nodes.

We need to allocate (ts_i, c_j) for each of the links. Equivalently, we need to colour the links, using the minimum possible number of colours: each colour corresponds to a (ts_i, c_j) tuple. For this, we first note that all the hop-1 links interfere with one another. This is because all of the antennas/radios at the landline location are going to be mounted on the same tower and will be in the vicinity of one another. So we have to allocate different colours to each hop-1 link. This then puts a lower-bound on the number of colours *necessary*. We will now show that this many colours are *sufficient* as well.

Now, we can take advantage of the tree structure, and the fact that traffic is from/to the landline, as follows. Denote the hop-1 nodes as N_i , and their respective links to the root as L_i . For a hop-1 node N_i , denote its set of hop-2 links as S_i . We now claim that for all the links in S_i , we only need one (ts_i, c_j) allocated. This is because the link L_i has been allocated only one slot. And for any uplink (downlink) traffic from (to) S_i from (to) the landline, L_i is the bottleneck.

Now, how do we colour the hop-2 links in each S_i ? For this, consider the example in Fig. 3. Here, S_1 consists of links between N_1 and P_i . Take the case of the link $N_1 - P_1$ using antennas B_1, B_2 . In the figure, L_2 is another hop-1 link using antennas A_1, A_2 . Now, consider the four antennas A_1, A_2, B_1, B_2 . Between any pair of antennas (A_i, B_j) , we see that not only are they separated by a long-distance, they have side/back lobes toward each other. More importantly, the tower/mast heights at N_1, N_2 , and P_1 are likely to be much lower than the tower height at R (see Sec. 2). This means that N_1 & N_2 , or P_1 & N_2 are very unlikely to have line-of-sight (due to obstructions and the earth’s curvature). Due to these factors, L_2 will be non-interfering with the link $N_1 - P_1$. Similarly, L_2 will be non-interfering with all the other links in S_1 too.

Generalizing the above example, we can intuitively see that for each S_i , several non-interfering L_j will exist in a large network: for a given S_i , all links L_j in the “other” side of R as L_i will likely be non-interfering. And S_i can be assigned the same colour as a non-interfering L_j .

To colour each S_i , we can now consider the problem of bipartite perfect matching: for each S_i , choose a non-interfering L_j (and allocate the same colour for S_i as L_j). A perfect matching will very likely exist, since each S_i has several possible L_j . Applying known polynomial algorithms for bipartite perfect matching will hence give us an optimal (ts_i, c_j) allocation in the long-distance network.

We have implicitly assumed above (Fig. 3) that L_1 and L_2 are al-




Figure 3: Non-interfering S_1 and L_2

located only one (ts_i, c_j) slot. In general, different links may need be assigned different number of slots, based on traffic requirements. But this can easily be accommodated. Suppose we count traffic requirement in units of b $Kbps$. Then if L_i has a traffic requirement of k units, simply consider it as k different links, and the above approach would still work⁵.

In the above approach, we have essentially identified a set of links which are mutual interfering (hop-1 links). These links certainly require different colours. Then we have argued that the same set of colours can be used for the hop-2 links too. Such an approach is possible essentially because there is a structure to the spatial reuse pattern.

Related to the approach outlined above, there are many open issues. First, extending the above approach to trees of depth three will be of value. Much like above, we could seek to colour the hop-3 links too using the same set of colours as those used for the hop-1 links. But we doubt if a generic consideration of arbitrary depth will be of much practical value.

Next, in the operation of the hop-1 links, we have not considered SynTx, SynRx as in 2P [10]. On the one hand, achieving SynTx/SynRx with sector antennas will be a significant engineering challenge than with directional antennas, since the side/back lobe patterns are more prominent for the former. On the other hand, incorporating 2P in the above formulation is a possibility which requires further exploration.

Time-slot, channel allocation in LACNs

Given the above approach to (ts_i, c_j) allocation in the long-distance network, how may we go about the same in each local-access network? Our approach is based on the following insight. Observation **O3**: for each LACN, the long-distance link at its local-gateway is a bottleneck. The idea is to use **O3** to show that there is enough slack for scheduling within each LACN, and a simple scheme will suffice.

To see why **O3** is true, denote the local-gateway at $LACN_i$, as LG_i . And denote the total capacity in one channel of operation is C . For instance, for 802.11b, C would be 11 $Mbps$ at the PHY layer (about 6 $Mbps$ at the MAC layer). The total capacity using k channels is thus kC ($k = 3$ for 802.11b/g). Also denote the total traffic to/from $LACN_i$, via LG_i , as C_i .

O3 is equivalent to saying that $C_i \ll C$. To see this, let the total number of nodes in the LDN be T . So there are T LACNs. Suppose the traffic requirements at these T LACNs are uniform,

⁵A perfect matching would still exist unless the traffic requirements are highly imbalanced.

then $C_i = kC/T$ at most. Typically, T is large for a network, and k is small. Thus we have **O3**: $C_i = kC/T \ll C$. This would be true unless the traffic requirements at the LACNs are highly skewed.

We next observe **O4**: we will have at most two channels occupied at LG_i . This is also easy to see. Suppose that LG_i is a hop-1 node, and S_i is the set of hop-2 links at that node. Denote the allocation to the hop-1 link at LG_i as (ts_{i1}, c_{j1}) , and the allocation to S_i as (ts_{i2}, c_{j2}) . Thus only two channels c_{j1} & c_{j2} are used at LG_i , and **O4** holds. Similarly, it holds for hop-2 nodes also. This implies that, under 802.11b/g, we have at least one channel entirely free for operating $LACN_i$.

We are now ready to present our solution for scheduling in $LACN_i$, using **O3** and **O4**. Fig. 4 illustrates how up to T/k hops can be supported easily. The figure assumes that all traffic is in the downlink (from landline). When LG_i gets a packet destined to a node D , it determines a route to D . Since the link-abstraction is valid, just using a simple minimum-hop scheme to determine such a route would work. It then sends the packet to D , operating each hop along the route one after another. In such operation, the time taken to send B bytes along h hops would be: $h \times B/C$, since each hop's (single-channel) capacity is C . Now, the time taken for B bytes to arrive over the LDN at LG_i is $B/C_i = T/k \times B/C$. Hence there will be no queue build-up at LG_i so long as $h < T/k$.




Figure 4: Scheduling in $LACN_i$

A similar argument applies in the uplink direction too (D to landline), and for any mix of downlink/uplink traffic. Now, considering a network of reasonable size of say, $T = 30$ nodes. And assuming 802.11b/g operation using three channels, we can easily support $T/3 = 10$ hops in each $LACN_i$ under the above scheme. Now, typical villages are within a km in length/breadth. Hence we expect $LACN_i$ to be only 3-4 hops in practice. The above approach can thus support most practical scenarios.

In the approach described above, there is a subtle but important challenge, not encountered in prior multi-hop TDMA schemes in the literature. In $LACN_i$, the schedule to be implemented changes depending on the destination D (or source, for uplink traffic)! The schedule has to be determined and changed at a fine granularity. But we believe that this is possible; this is one of the topics in the next section.

4. TDMA MAC: IMPLEMENTATION ISSUES

There are three main challenges we face in the above design of our TDMA scheme. (1) Like all multi-hop TDMA schemes, we have the significant issue of achieving time synchronization, in a potentially large network. (2) In FRACTEL, we expect to have changing traffic patterns, and at any time only a subset of the nodes in the network may be active. So it does not make sense to have any static allocation. This then raises the challenge of deciding the granularity at which TDMA allocation should be done. (3) Finally,

as mentioned in Sec. 3 above, in each $LACN_i$ we need to perform scheduling at a fine granularity, depending on which node in $LACN_i$ is the source (destination) of traffic to (from) the landline.

To address the above challenges, our proposed approach consists of four strategies: (a) use the *hierarchical* structure of the network, (b) use *centralized* algorithms for synchronization and scheduling, (c) use a *multi-hop connection-oriented* link layer, and (d) use fine granularity scheduling in each LACN. These four strategies fit in nicely with one another as we explain below.

Recall from Sec. 2 that each LACN is independent of other LACNs. So ideally we would like to have independent synchronization mechanisms in each LACN, and an independent synchronization mechanism in the LDN. But then, each LACN's local-gateway has long-distance links which interfere with the LACN's links. Given this, how do we avoid synchronizing the entire network?

O4 in Sec. 3 answers this: we can have an entire channel of operation for a LACN. Thus there is no need to synchronize each LACN with the LDN, or the LACNs with each other.

Now, in the uplink direction (D to landline), how does LG_i know when to schedule for D ? To address this, we propose a *multi-hop connection-oriented* link layer. That is, before any traffic flows from/to D , it has to setup a connection with the landline. This enables us to categorize traffic flows as voice, video, and best-effort, much like in other connection-oriented approaches like 802.16 [13]. The connection has two distinct parts, in the LACN and in the LDN respectively. That is, both LG_i and the *landline* are aware of the connection.

Such connection categorization can then be used to schedule uplink (and downlink) traffic at LG_i . For instance, if a voice-flow has been setup, the allocation of uplink slots can simply be done periodically. For variable bit-rate video flows, we can augment periodic slot allocation with additional slots, as per the uplink queue at D . Like in 802.16, D can convey its requirement for additional slots by either piggybacking such requests with its data packets, or in periodic polling packets [13].

Now, to implement such a scheme, LG_i repeatedly schedules a *multi-hop downlink* sub-frame, and a *multi-hop uplink* sub-frame. This is similar to 802.16 framing, albeit with the significant distinction that it is multi-hop. Achieving such multi-hop uplink/downlink framing is an interesting avenue for further exploration. But we believe that it is possible, for the following reason. In Sec. 3, we explained that it is possible to support T/k hops in each $LACN_i$. Suppose we restrict it to allow only $T/2k$ hops (5 hops for $T = 30$ and $k = 3$), then we straight-away have as high as a 100% leeway for any overheads in implementing multi-hop downlink/uplink framing. And in practice, supporting just 4-5 hops would be quite adequate.

Another point to note in this context is related to link-layer ARQ. Recall that we can have the link abstraction in FRACTEL, and achieve close to 0% error rates on individual links. This implies that it would be more efficient to have link-layer ACKs over multiple hops (between D and LG_i). This too fits in well with our multi-hop frame based mechanism: the ends which are involved in the framing are also the ones involved in any ARQ retransmissions.

The above scheme thus allows for independent synchronization and scheduling within each $LACN_i$. And independence between the LDN and each $LACN_i$ too. For the LDN, we can have centralized control at the landline. And for each $LACN_i$, we can have centralized control at their respective LG_i . In general, centralized design is shunned upon in networks, due to the poor fault-tolerance. But in FRACTEL, it is alright to have the landline or the LACN local-gateway nodes have centralized functionality. This is because if LG_i is down, in any case $LACN_i$ cannot send or receive traffic.

Similarly, in the LDN, if the landline is down, the entire network is down anyway.

Scaling too is not affected due to centralization of functionality since we use it along with the notion of hierarchy: each LG_i is responsible only for its own $LACN_i$, and the *landline* is responsible only for the LDN. The combination of hierarchy and centralized control thus achieves scaling.

The connection-oriented approach will also help in appropriate scheduling at the landline, for the LDN. During connection setup, the LDN is aware of flow setup, and how much traffic the flow will generate. This will help us address the challenge of dynamic scheduling to adapt to the changing traffic pattern.

While we have outlined a possible approach, clearly, several issues need in-depth exploration. How can we achieve multi-hop downlink/uplink framing? What are going to be the overheads involved? How do these overheads vary when we seek to achieve lower delay (smaller frame duration)? How exactly can the scheduling be done for each category of traffic?

Another implementation issue of significance is going to be the following. How exactly may we achieve multi-hop synchronization using off-the-shelf 802.11 hardware? The work in [14] and [15] have shown prototypes of TDMA-based single-hop MAC implementations based on off-the-shelf 802.11 hardware. The work in [15] has shown how we can implement an 802.16-like TDMA MAC protocol over "cheap" 802.11 cards. Questions still remain of how we may tap into the existing (single-hop) time synchronization function of 802.11 to implement multi-hop synchronization, and how efficient such an approach can be.

An issue of considerable importance is related to dynamic channel & time-slot scheduling. Clearly, in a large network, we do not want to change the channel or time-slot allocation completely each time a new connection is formed. It would be a significant challenge to achieve dynamic scheduling with minimal disruption to an existing allocation or schedule.

5. CONCLUSIONS

In this paper, we have presented FRACTEL, a mesh network for deployment in rural settings. We argue that FRACTEL is starkly different from generic (urban) mesh networks which have been considered in prior literature. In FRACTEL, individual links can be built for good, predictable performance. To support QoS for real-time applications, we envision a TDMA-based MAC. Several unique aspects in a FRACTEL network make us consider technical issues in this context with a fresh perspective. We have proposed novel approaches to TDMA scheduling and channel allocation in FRACTEL, and have articulated various issues which warrant in-depth consideration.

Although our discussion has been centered around 802.11-based mesh networks, many of our approaches are more generic. For instance, the time-slot and channel allocation mechanisms we have discussed could also be applied for 802.16 mesh networks under similar settings.

Acknowledgment

We thank Sayandeep Sen, Dattatraya Gokhale, Zahir Koradia, Pradeep Gopaluni, and Gaurav Chhawchharia for the various brainstorming discussions during the early phases of the FRACTEL project. We are grateful to the anonymous reviewers for their comments on this paper.

6. REFERENCES

- [1] IEEE P802.11, The Working Group for Wireless LANs.
<http://grouper.ieee.org/groups/802/11/>.
- [2] Pravin Bhagwat, Bhaskaran Raman, and Dheeraj Sanghi. Turning 802.11 Inside-Out. In *HotNets-II*, Nov 2003.
- [3] Rabin Patra, Sergiu Nedeveschi, Sonesh Surana, Anmol Sheth, Lakshminarayanan Subramanian, and Eric Brewer. WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks. In *USENIX NSDI*, Apr 2007.
- [4] Bhaskaran Raman and Kameswari Chebrolu. Experiences in using WiFi for Rural Internet in India. *IEEE Comm. Magazine*, Jan 2007.
- [5] Kameswari Chebrolu, Bhaskaran Raman, and Sayandeep Sen. Long-Distance 802.11b Links: Performance Measurements and Experience. In *MOBICOM*, Sep 2006.
- [6] Daniel Aguayo, John Bicket, Sanjit Biswas, Glenn Judd, and Robert Morris. Link-level Measurements from an 802.11b Mesh Network. In *SIGCOMM*, Aug 2004.
- [7] Mansoor Alicherry, Randeep Bhatia, and Li (Erran) Li. Joint Channel Assignment and Routing for Throughput Optimization in Multi-radio Wireless Mesh Networks. In *MOBICOM*, Aug/Sep 2005.
- [8] Murali Kodialam and Thyaga Nandagopal. Characterizing Achievable Rates in Multi-hop Wireless Networks: The Joint Routing and Scheduling Problem. In *MOBICOM*, Sep 2003.
- [9] Ashish Raniwala, Kartik Gopalan, and Tzi-Cker Chiueh. Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks. *ACM MC2R*, 8(2), 2004.
- [10] Bhaskaran Raman and Kameswari Chebrolu. Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks. In *MOBICOM*, Aug/Sep 2005.
- [11] Sayandeep Sen and Bhaskaran Raman. Long Distance Wireless Mesh Network Planning: Problem Formulation and Solution. In *WWW*, May 2007.
- [12] Bhaskar Ramamurthi. Broadband Wireless Technology for Rural India. In *URSI-GA*, Oct 2005.
- [13] IEEE 802.16 WirelessMAN.
<http://www.ieee802.org/16/>.
- [14] Michael Neufeld, Jeff Fifield, Christian Doerr, Anmol Sheth, and Dirk Grunwald. SoftMAC - Flexible Wireless Research Platform. In *HotNets-IV*, Nov 2005.
- [15] Pavan Kumar. Design, Implementation, and Evaluation of new MAC Protocols for Long Distance 802.11 Networks. Master's thesis, IIT-Kanpur, May 2006. *with Bhaskaran Raman*.