

Lo³: Low-cost, Low-power, Local Voice and Messaging for Developing Regions

Bhaskaran Raman
Dept. of CSE,
IIT Bombay
br@cse.iitb.ac.in

Kameswari Chebrolu
Dept. of CSE,
IIT Bombay
chebrolu@cse.iitb.ac.in

ABSTRACT

System design for deployment in developing regions, especially for rural areas, must consider two metrics of prime importance: system cost and power consumption. In this work, we propose a novel approach for building a local communication system, within a village. Our approach is based on IEEE 802.15.4 as the underlying technology. Although designed for a completely different application space (embedded wireless sensing), and although the radio capacity is small, we argue that 802.15.4 is ideally suited to minimize cost as well as power consumption. It is possible to enable many useful voice-based applications. The use of the radio in this non-traditional setting brings up several technical challenges however. This paper articulates these research issues.

1. INTRODUCTION

The design of communication systems for rural deployment in developing regions of the world presents unique challenges. The primary optimization criteria here are cost and power consumption [1]. Every aspect of the system, including the infrastructure (not just the end-user device) must consider these two criteria. Performance optimization is an important, but secondary criterion. In this paper, we present Lo³, a **low-cost, low-power, local** communication system for use in village settings in developing regions¹.

In Lo³, we wish to enable applications such as two-way interactive voice, stored voice/text messaging, community alerts or warnings, community radio, etc. Such a system can be viewed as a PABX (private automatic branch exchange), for use within a village. Voice-centered applications are especially suitable for rural areas in developing regions, where the literacy levels are low [2]. While traditional web applications require relevant content creation to promote usage, voice applications can see immediate usage.

The overall architecture of Lo³ is based on IEEE 802.15.4 [3].

¹Lo³ is pronounced as “loh-cube”.

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.

Copyright 2009 ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

This suits us well, given the goals of low cost and low power consumption. We argue that the use of a mesh network helps in cost reduction; the infrastructure cost of such a network can be as little as about U.S.\$1,500. This is much less than alternate possibilities such as cellular systems (GSM/CDMA).

For effective power-savings, we propose that the network nodes be duty-cycled whenever there is no traffic through them. They periodically sleep and stay awake. We thus have a duty-cycling, mostly sleeping mesh network, the relevant parts of which awaken when necessary to support communication.

802.15.4 was originally designed for embedded wireless sensing². The radio is popular in many recent wireless sensor networking (WSN) applications, such as Volcano monitoring [4], BriMon [5], etc. Such WSN applications focus on low cost and low power consumption, but do not consider support for real-time applications. To our knowledge, the only system to have considered real-time voice on 802.15.4 is FireFly [6]. But unlike Lo³, FireFly does not consider power optimization through duty-cycling. So our system is unique in terms of its combination of three considerations: cost minimization, power savings (duty cycling), and real-time application support. This presents many novel technical issues in terms of hardware design, MAC, routing, and mobility support, which are unique to our setting. These issues require careful consideration, as we articulate in the paper.

In the next section (Sec. 2), we detail the motivation behind Lo³. We then describe the various technical issues in this setting, in Sec. 3. Sec. 4 presents a few points of discussion, and Sec. 5 concludes the paper.

2. MOTIVATION AND GOAL

2.1 Communication Technology Choices

What communication technology choices can be considered for rural deployment? What applications can they support? Considering these questions, we now motivate Lo³.

Since wiring is expensive, the use of wireless is an obvious choice. *Cellular technology* (GSM/CDMA) is quite mature for voice applications. But its primary shortcoming is that base station cost can be as high as \$10-100K, depending on the configuration. Furthermore, base station power consumption is also high (0.5-1 kW), requiring huge power backup systems and also cooling systems. That cellular technology is not appropriate is also quite evident from the fact that in countries

²The standard is often used in conjunction with ZigBee.

like India, deployment in rural areas is poor (except near national highways), despite dense deployments in urban areas.

It is to address this shortcoming that *WiFi-based* systems have been suggested for rural settings [1]. WiFi offers data rates in the Mbps range (1-54), and can potentially support broadband connectivity, VoIP, and real-time video conferencing [7]. While the cost of the radio itself is quite low ($< \$10$), the infrastructure cost of towers in long-distance WiFi networks are known to be substantial [8, 9]. And in a system seeking to provide video-conferencing, there are other significant costs too; the Ashwini project reported a per-site cost of \$1.5K each for (a) the electrical fixtures required in a video-conferencing room, and (b) the power backup costs [10]. Overall, the Ashwini WiFi mesh network and the associated applications had a setup cost of about \$0.5 million for 30+ villages, with plans to recover the amount (break-even) over a period of five years.

Clearly, such a scale of operation with a break-even period of several years may not be attractive in all settings. Given this, we ask: how much can we scale down the cost and power requirements, while still providing a communication service which is useful? It is this thought process which has led us to consider the following goal:

Lo³ goal: design a low-cost, low-power, local voice communication system.

More explicitly, in contrast to the WiFi systems deployed in [10, 7], we scale down the application set to just local voice applications and seek a significant trade-off in terms of system cost and power consumption.

802.15.4 radio technology is inexpensive as well as power efficient. In 802.15.4-based platforms, the power consumed can be as low as 60-360mW based on the transmit power setting: 0-20dBm (the Tmote platform consumes 60mW at 0dBm transmission). Importantly, these platforms also have good support for duty-cycling; the sleep power consumption can be as low as a few micro-Watts. In comparison, WiFi access points consume as much as a few Watts, and have poor support for duty-cycling. (In terms of cost, WiFi and 802.15.4 radios are comparable). We thus choose 802.15.4 for Lo³.

We originally also considered *Bluetooth* as a possible choice as it offers good data rates (1Mbps). But Bluetooth is more power hungry and costlier when compared to 802.15.4 [11]. Further, off-the-shelf devices do not have support for any multi-hop scatternet, and the MAC protocol does not lend itself to easy tinkering using off-the shelf equipment. Bluetooth link establishment delay is a concern too: it can take as long as 10sec to establish a *single* link.

In this paper, we consider 802.15.4 operation in the 2.4GHz band only, although it can operate in other bands too. This is primarily because the 2.4GHz band is license-free globally.

While 802.15.4 is the most appropriate for Lo³, the low data rate (250Kbps) and duty cycled operation of the hardware raise many interesting challenges as we describe in Sec. 3.

2.2 Design Requirements

We now elaborate on some important aspects of the above-stated goal.

In Lo³, by “local” communication, we mean to serve users within a village area of about a few kilometres in diameter. For instance, the typical expanse of a village in rural regions in India is an area of about 2-3km diameter. Houses are typically within a 1km² area or less, while there are farms in the

vicinity³. While carpet coverage is desirable, we believe that even *hotspot coverage*, at important buildings and locations in a village, can be quite useful.

Cost is an important consideration since the paying capacity of users is limited. Cost must be minimized for both the infrastructure as well as the end-user device (handset). Power consumption too is important to minimize, not only for end-user devices, but also for infrastructure devices. This is because there is a lack of reliable power supply in villages in rural areas of developing regions [1, 2]. Power-off-the-grid may often not be available for several days at a stretch, the use of batteries is hence inevitable. The use of solar power is possible, but not without drawbacks. Solar panels add cost to the system, and they are also an attractive target for theft (due to the necessarily conspicuous installation, the inconvenience/impossibility of placing them under lock-and-key, and their ready resale value). Thus it is not wise for system design to bank on solar power; we target our system to be able to operate out of battery power for several weeks to months, requiring maintenance in terms of recharge only very rarely.

In terms of performance requirements, our system should support good voice quality. In terms of usage volume, say that we have a user population of a 100-200, among a total village population of 1000+ people. While it is desirable to support as many simultaneous calls as possible, we believe that to begin with, support for even as few as 4-5 simultaneous calls would be good enough, so long as the system is affordable. This is because we are targeting Lo³ for regions which have little or no communication support today. As system usage grows, it is desirable that the system allow for higher capacity, with the potential addition of more infrastructure.

3. LO³ DESIGN, TECHNICAL CHALLENGES

We now present various aspects of Lo³ design which we have worked out, and also articulate the various unresolved technical issues.

3.1 Choice of Topology: Star versus Mesh

The first design aspect is that of the network topology. We consider two possibilities: star versus mesh topology. A star topology essentially means a tall central tower that hosts an infrastructure node to cater to end user handsets in the village. All communication is via this central infrastructure node. The advantage of this approach is the simplicity; it makes protocol design easy. However it has many disadvantages.

First, as mentioned earlier, the tower cost can be quite high, especially given that it has to cover a few kilometres. A tower height of at least 25-30m would be required to clear trees and other obstructions. Such a tower itself would cost several thousands of dollars [9]. The cost also would have to be paid upfront. And importantly, even with a tall tower, there are likely to be many shadow regions (behind trees/buildings), especially in the 2.4GHz non-penetrative band. Relays would then be needed for such regions, necessitating deviation from the star topology.

In the alternative mesh topology architecture, communication between two end devices may be multi-hop via one or more infrastructure nodes. The infrastructure nodes may be deployed on rooftops (5-10m height). At these heights, we are

³This is in the plain regions, not in hills.

likely to get a range of several hundreds of metres. For instance, 802.15.4 measurements in [5] reported a range of over 400m using 0dBm transmissions, and 8dBi omni-directional antennas mounted at a height of less than 3m. Mounting at a higher height (5-10m) will only improve the range. Given this, an area of about 3km diameter can be (hotspot) covered by a mesh of about 20-30 infrastructure nodes. Also, the use of a mesh allows incremental deployment.

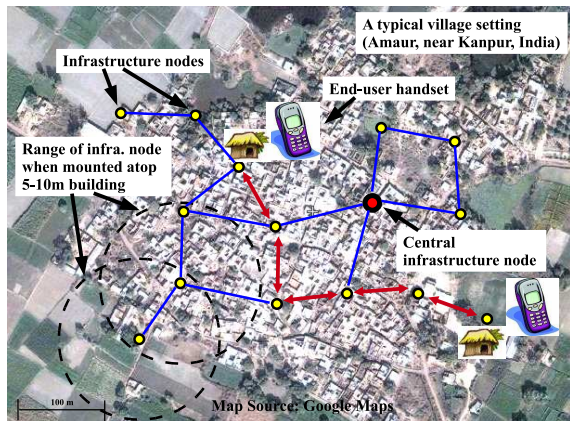


Figure 1: Lo³ architecture and usage model

Fig. 1 depicts the Lo³ mesh architecture and usage model. A mesh topology however gives rise to challenges in terms of protocol design, especially if we want power optimization through duty-cycling. But we believe that these challenges can be met, as we explain below.

3.2 Choice of Hardware

There are two categories of hardware that are needed in Lo³: the end user handset, and the infrastructure node. The hardware design needs to be cost and power efficient.

The handset, while being inexpensive and power-efficient, must at the same time be powerful enough to handle good audio encoding/decoding. This is because raw audio, or simple encodings like PCM (64Kbps), are too wasteful given the limited radio capacity (250Kbps). Encoding/decoding is a challenge with most low end processors since the encoding/decoding task has to finish within a small interval (e.g. within 20ms for GSM).

We initially considered two off-the-shelf platforms for the handset: the MSP430-based Tmote platform, and the Jennic platform (www.jennic.com) based on a 32MHz, 16-bit processor. However, both these possibilities present problems. The MSP430 processor is not powerful enough for efficient codecs; and the performance of a codec such as 16 Kbps AD-PCM is not very good [6]. While the Jennic platform uses a more powerful processor, it unfortunately does not allow enough flexibility to implement our own network stack: it comes with a default ZigBee/802.15.4 stack.

We are currently exploring the design of the handset using TI's low-power DSP chips from the TMS320C55x series. These chips provide enough CPU power and memory to implement a codec such as the GSM codec. They are inexpensive (about \$7) and also power efficient: as little as 150mW when active and only a few tens of micro-Watts when asleep.

For the infrastructure node, we may be able to use an MSP430-

based platform, like the Tmote. One flexibility which would be useful in infrastructure nodes is a dual-radio configuration. Mesh networks based on dual-radio nodes offer higher capacity, by enabling better channel reuse (802.15.4 has 16 independent channels of operation in the 2.4GHz band). We are unaware of any hardware platform that supports dual 802.15.4 radios. The challenge in designing such a platform is that the two radios should not interfere with each other due to close proximity. We also need to ensure that the interface between the two radios does not turn out to be a bottleneck.

Whether it would suffice to use a low-end processor such as MSP430 for most infrastructure nodes (single radio or dual radio) requires further validation.

In summary, there is no off-the-shelf platform which we can readily use for Lo³. We have to design the required components, with due consideration for cost and power.

3.3 Duty-Cycled Operation

A common approach to conserve power is duty-cycling, where nodes sleep and stay awake periodically. This would be necessary to achieve an infrastructure up-time of several weeks or more without battery recharging. Since we expect a low call volume, we can have the network nodes duty-cycle when there is no traffic on the network. Even during a phone call, only the relevant nodes (on the path between the two communicating nodes) need to be fully awake; the other nodes can continue to duty-cycle.

Such duty-cycling can however pose a challenge to any protocol since an intended message destination may not be awake when required. This often results in a tradeoff between latency and power savings. By latency, here we mean the time taken to fully awaken the relevant nodes for a particular phone call, i.e. call setup time. Larger the sleep time, greater the power savings, but larger the expected call setup time. We note here that it is alright to have a call setup time of 5-10sec without seriously affecting the system usability, since users are anyway used to waiting a while after dialing.

There are two main approaches possible for duty cycling. One is a time synchronized approach where all the nodes synchronize their clocks and sleep and wakeup roughly at the same time. If one node wants a set of other nodes to remain awake for a phone call, it sends a message during the wakeup time to the other nodes saying so. Once the data exchange, or phone call, is done, the nodes can go back to sleep.

An alternative possibility is to have nodes duty cycle independent of one another. A node can wake up a neighbouring node asynchronously via continuous probes, one of which is received by the destination during its awake interval (similar to X-MAC [12]). And this process can be repeated one hop after another. This does not need any time synchronization.

The primary issue with the time-synchronized approach is that multi-hop time synchronization is considered to be hard in practice. However, we believe it is practical to achieve such synchronization with manageable synchronization error. From [5], we know that it is possible to achieve a synchronization error of only about 30-60 μ s per hop, using a simple flooding synchronization protocol. For the network size we expect, suppose we have 10 hops, this means a maximum synchronization error of under a milli-second, which appears manageable.

However, there are several open issues in this regard. First, is it indeed possible to do duty-cycling stably (without going

out of sync) even in the presence of such synchronization errors? From another perspective, is it possible to sustain the low level of errors over time, even though the network is duty-cycling? This is significant since any re-synchronization can happen only when nodes are awake. It is worth noting in this context that time-synchronization protocols in literature (e.g. see FTSP [13]) typically assume that nodes are awake all the time. In our setting, how often should we re-synchronize? And finally, the question of primary interest is: what is the effective duty-cycling ratio we can achieve?

A time-synchronized network also brings benefits for other network functionalities, such as the MAC, as we explain below.

3.4 Medium Access Control (MAC)

The issue of MAC comes up in Lo³ specifically among the nodes which are awake (not duty-cycling) for the ongoing (one or more) phone calls. Though 802.15.4 specifies a MAC protocol, the standard does not address many important issues in multi-hop networks. So we consider the issue from scratch.

Given the scarcity of bandwidth associated with 802.15.4 (250Kbps) and the need to support delay sensitive voice application, it would be useful to utilize all the available channels (16) to achieve channel reuse. There are three design decisions to be made for the MAC in a multi-channel mesh.

CSMA or TDMA? It is well known that a CSMA/CA MAC will perform poorly in a multi-hop mesh topology, given the scope for self-interference (i.e. among adjacent links on the same path). The contention will give rise to large delay and jitter, while hidden node situations will cause packet losses.

A TDMA MAC has the potential to enable low delay and jitter, required for voice communication. Furthermore, it can provide good support for duty-cycling and power savings in infrastructure as well as mobile nodes. TDMA requires multi-hop time synchronization, but as mentioned earlier, this appears possible for the envisioned scale of operation.

Centralized or distributed MAC? A centralized MAC makes global coordination (specifically for TDMA) much simpler and the resulting MAC simple. This can be very useful on a platform with limited processing capabilities such as the MSP430-based Tmote. This also makes testing and deployment easy. Centralized optimization of aspects such as time-slot scheduling, channel allocation, flow admission control, nodes' battery power, etc. are also likely to be simpler than their distributed counterparts. Such optimization is especially important given the limited radio capacity.

Centralized control in general has the disadvantages of lack of scalability and fault tolerance. On the issue of scaling, we believe that a careful design can provide bottleneck-free operation for the envisioned network sizes. The issue of fault tolerance can be addressed by having backup nodes for the centralized functionality. Or a deployment may even choose to ignore the issue of fault tolerance initially: some availability is better than the zero availability today.

Connection-oriented or connection-less? A connection-oriented MAC means that the higher layer can use it only after a connection formation phase. This fits in well with a TDMA-based multi-channel approach: the connection formation phase can be used to specify the time-slot and channel of operation for each node in the data path.

So our design choice is to have a centrally controlled, connection-oriented TDMA MAC. It is worth noting that while literature

is abound with TDMA scheduling approaches in mesh networks, there is no multi-hop TDMA-based wireless system deployed to-date, to our knowledge. There are several open issues in this regard which need to be addressed. Is the synchronization error indeed manageable? What performance hit does it entail? What is the overhead to maintain synchronization in presence of ongoing voice calls? How exactly does the TDMA MAC handle dissemination of any computed schedule? Can the MAC indeed be implemented on low-end hardware platforms? How many calls can be supported? How can one implement admission control which is necessary given the limited capacity?

Only a careful design, and evaluation using a prototype can reveal subtleties, validate the design choices, and answer the above questions.

3.5 Routing Protocol

We now consider the issue of routing in the Lo³ infrastructure network.

We first discuss the choice of a routing metric. Prior measurement work using 802.15.4 radios has revealed that it is possible to achieve close to 0% link error rates, when operating (static) links at RSSI values above a certain threshold [14]. So with careful deployment, we expect the infrastructure links to be stable, i.e. no varying error rates with time. So routing based on metrics like ETX/ETT (expected transmission count/time) [15] etc. are not necessary in our setting.

Since we are operating with multiple channels, it is possible to route data based on channel diversity of a given path. But with 16 independent channels, there will likely be sufficient diversity in the network. So distinguishing one path from another using a metric like WCETT [15] is also unlikely to be useful.

A metric that seems to be of interest (apart from the shortest hop count) is the residual battery life of the nodes. Here too, since we expect the network lifetime to be of the order of several weeks to months, updating this metric can happen at a very coarse time scale: minutes/hours.

An important question in a duty-cycled network, is *when* the routing protocol should run. The above discussion on the routing metric helps answer this. Given the coarse granularity of metric update, we can have the routing protocol run only infrequently, and cache the routes. Any node failure (or recovery) can be detected, for instance during a call establishment, or in the middle of a voice conversation. And this can trigger the routing protocol to run again. This "pause-resume" approach to routing also fits in well with the requirement of route pinning, for real-time applications.

One subtle requirement of a routing protocol intended to run only infrequently is that nodes need to know *when* routing activity can be *paused*, so that the nodes can go to sleep. This is very much unlike traditional routing where the protocol runs continuously, and there is no notion of "pausing". Current distributed protocols do not tell when the routes have stabilized. And designing a distributed protocol that can tell this precisely (or even roughly) does not seem very straightforward. This is yet another reason to have centralized protocol control.

3.6 Mobility Support

We wish to support portability and mobility in our setting. While support for vehicular speed is desirable, pedestrian mo-

bility is more important. We need to handle two aspects: paging and handoff.

Paging: This is the process of finding out where in the network the mobile node is. One possible way to achieve this a “network-pull” model where the network floods a query to locate a specific mobile. However this can be quite inefficient, both in terms of power consumption and latency. Furthermore, such flooding must also take care not to disturb other on-going calls; an on-going call may have scheduled time-slots for its packet transmissions, the flooded paging query must not use these.

An alternate solution is to use a “mobile-push” model where the handset registers its location. When exactly this registration should happen is a significant question. Cellular network literature is abound with various approaches to this. But there are additional subtleties in Lo³. In the cellular infrastructure only the mobiles duty-cycle, while the infrastructure is all-powerful. However, in Lo³, the infrastructure itself duty-cycles.

Handoff: Supporting handoff in a duty-cycled environment seems quite tricky. Typically, handoff in mobile networks is supported by keeping the neighbouring nodes around a mobile handset ready, in anticipation of an impending handoff. Doing this in Lo³ means that potentially several neighbouring infrastructure nodes must be kept awake unnecessarily. This is wasteful of power.

Further, low-latency handoff potentially requires duplicating the voice call data packets on multiple paths near the mobile (so that the mobile continues to receive packets even if it moves in range of a neighbouring infrastructure node). This leads to wastage of the already constrained bandwidth. However, if we do not do such duplication, we need to wakeup the right infrastructure node on demand and setup the new path. This means higher latency that may affect voice call quality during the handoff.

The tradeoffs involved in paging and handoff mechanisms require careful evaluation. We are currently in the process of building a simulator for the initial exploration.

3.7 Applications

Apart from regular two-way voice, Lo³ can be used for other applications too.

Stored Voice Messaging: It would be useful to enable stored voice messaging for cases where the intended call destination is switched off or the call could not be admitted. In this case, the message can be stored and delivered in leisure to the other end.

Broadcast: Some applications like community alert, community radio etc. will need a broadcast tree to send packets to all the end-users. The computation of the broadcast tree can be part of the routing protocol. However care must be taken that a broadcast does not disrupt ongoing voice calls.

4. DISCUSSION

Lo³ prototype: We are in the process of building a Lo³ prototype. As of this writing, our prototype includes: (1) multi-hop time-synchronization, (2) multi-channel TDMA operation, (3) a connection-oriented MAC. The prototype is based on TinyOS, and runs on the Tmote sky 802.15.4-compliant platform. This *does not* include the end-user handset functionality, although we have been able to achieve real-time audio

streaming between two laptops, via a 802.15.4 path. While it is beyond the scope of this paper to detail the prototype, it is worth mentioning that we have been able to achieve 60+ Kbps over a 9-hop linear path, with time-slotted, multi-channel operation. More importantly, we have been able to account for all of the overheads involved. This has given us confidence that a multi-hop TDMA-based approach is indeed feasible.

Lo³ simulator: In parallel to the prototype, we have also built a simulator to quickly study various MAC-related design issues, as outlined in Sec. 3.4. Our initial findings indicate that duty-cycled operation over a network of a few tens of nodes is feasible without hitting scaling bottlenecks. We are in the process of further detailed evaluation.

Back-of-the-envelope cost, power, and capacity analysis: Only a complete prototype can tell for sure if we have met our design goals, but a quick calculation is encouraging.

System cost: The bill-of-materials for the end-user handset can be as little as \$30-40, including the DSP chip (\$7), an 802.15.4 radio (say \$6-8), audio mic/speaker (say \$6-8), and other miscellaneous components. For the infrastructure node too, the bill-of-materials of a platform similar to Tmote, but with just the necessary components (MSP430 processor, radio, serial, etc.) can be as low as \$30. So a total infrastructure node cost of about \$50-60, including batteries, antenna, packaging, etc. appears feasible. This means that a 20-25 node infrastructure will cost under \$1.5K.

Power consumption: Suppose that we have an aggregate call duration of 2 hours per day, and that we have 20% effective duty-cycling. Let us assume a node power consumption of 100mW when a node is active (the Tmote platform consumes about 60mW while transmitting or receiving at 0dBm). So the power consumption per day is: $100mW \times 2H + 20\% \times 100mW \times 22H = 640mWH$. Suppose we use a 4.5AH, 12V battery at each of the infrastructure nodes (such batteries are about a third the size of a car battery; they are bulky, but that is alright for a static infrastructure node). We can effectively operate for $4.5AH \times 12V/640mWH \simeq 84days$.

System capacity: A single 802.15.4 link can provide a maximum data rate of 250Kbps. Our prototype achieves 60Kbps. If we use the G.723.1 codec, we require 6.4Kbps x 2 directions = 12.8Kbps per call. So we can support about four simultaneous calls in the worst case. Of course, in the average case, different calls could be in different parts of the mesh network, and more simultaneous calls can be supported. We believe this is sufficient for reasonable usage within a village. If necessary, the capacity can be enhanced by going for a dual-radio setting.

Integrating Lo3 with the outside world: Though Lo³ itself focuses on communication just within a village, it can very well be interfaced with external communication. For instance, Lo³ may be interfaced with a PSTN line, or with an end-point of a WiFi long-distance link coming into the village [1]. One challenge in interfacing with WiFi would be the management of the channels of the long-distance link as well as the local links, in the limited 83.5 MHz of free spectrum in the 2.4 GHz ISM band. Appropriate channel allocation should help to some extent in this regard.

Advanced call services: In abstract functionality, Lo³ is like a distributed version of a PABX (private automatic branch exchange) used in many enterprises. Some advanced call services in PABX systems include: call waiting/hold, automatic call back, etc. It would be interesting to explore whether we

can implement such services in the Lo³ system without compromising on the cost or power consumption aspects.

5. CONCLUSION

Lo³ is an example of a system designed from the ground up for deployment in rural settings in developing regions. An overarching design principle is that we value cost and power optimization more than performance optimization. Much of the design choices in the system, such as the use of IEEE 802.15.4, a mesh network topology, relatively impoverished hardware, etc. are the direct result of this principle. Although 802.15.4 was designed for use in industrial embedded sensor networks, the cost and power efficiency make it attractive for our application. Our position is that the use of 802.15.4 in this fashion is the right design point for enabling local communication in village settings. Although only a real deployment can confirm its usefulness, we expect that the voice-centered applications would catch-on. We have built a Lo³ simulator as well as an initial prototype: we are currently using these for detailed performance evaluation.

6. REFERENCES

- [1] Pravin Bhagwat, Bhaskaran Raman, and Dheeraj Sanghi. Turning 802.11 Inside-Out. In *HotNets-II*, Nov 2003.
- [2] Eric Brewer, Michael Demmer, Bower Du, Kevin Fall, Melissa Ho, Matthew Kam, Sergiu Nedeveschi, Joyojeet Pal, Rabin Patra, and Sonesh Surana. The Case for Technology for Developing Regions. *IEEE Computer*, 38(6):25–38, June 2005.
- [3] IEEE 802.15 WPAN Task Group 4 (TG4). <http://www.ieee802.org/15/pub/TG4.html>.
- [4] Geoffrey Werner-Allen, Konrad Lorincz, Matt Welsh, Omar Marcillo, Jeff Johnson, Mario Ruiz, and Jonathan Lees. Deploying a Wireless Sensor Network on an Active Volcano. *IEEE Internet Computing*, Mar/Apr 2006.
- [5] Kameswari Chebrolu, Bhaskaran Raman, Nilesh Mishra, Phani Kumar Valiveti, and Raj Kumar. BriMon: A Sensor Network System for Railway Bridge Monitoring. In *MobiSys*, Jun 2008.
- [6] Rahul Mangharam, Anthony Rowe, Raj Rajkumar, and Ryohei Suzuki. Voice over Sensor Networks. In *RTSS'06*, Dec 2006.
- [7] 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.
- [8] Bhaskaran Raman and Kameswari Chebrolu. Experiences in using WiFi for Rural Internet in India. *IEEE Communications Magazine, Special Issue on New Directions In Networking Technologies In Emerging Economies*, Jan 2007.
- [9] Debmalya Panigrahi, Partha Dutta, Sharad Jaiswal, K. V. M. Naidu, and Rajeev Rastogi. Minimum Cost Topology Construction for Rural Wireless Mesh Networks. In *INFOCOM'08 Mini-Conference*, Apr 2008.
- [10] Project Ashwini. <http://www.csdms.in/conflux/>, Oct 2005.
- [11] Lakshman Krishnamurthy, Robert Adler, Phil Buonadonna, Jasmeet Chhabra, Mick Flanigan, Nandakishore Kushalnagar, Lama Nachman, and Mark Yarvis. Design and Deployment of Industrial Sensor Networks: Experiences from a Semiconductor Plant and the North Sea. In *SenSys*, Nov 2005.
- [12] Michael Buettner, Gary V. Yee, Eric Anderson, and Richard Han. X-MAC: A Short Preamble MAC Protocol for Duty-Cycled Wireless Sensor Networks. In *SenSys'06*, Nov 2006.
- [13] Miklos Maroti, Branislav Kusy, Gyula Simon, and Akos Ledeczki. The Flooding Time Synchronization Protocol. In *SenSys*, Nov 2004.
- [14] Bhaskaran Raman, Kameswari Chebrolu, Naveen Madabhushi, Dattatraya Y Gokhale, Phani K Valiveti, and Dheeraj Jain. Implications of Link Range and (In)Stability on Sensor Network Architecture. In *WiNTECH*, Sep 2006.
- [15] Richard Draves, Jitendra Padhye, and Brian Zill. Routing in Multi-Radio, Multi-Hop Wireless Mesh Networks. In *MOBICOM*, Sep 2004.