

# Deploying a Rural Wireless Telemedicine System: Experiences in Sustainability

*Sonesh Surana, Rabin Patra, Sergiu Nedevschi, and Eric Brewer*  
University of California, Berkeley

**A novel ICT project in rural India uses long-distance Wi-Fi networking to enable high-quality videoconferencing between eye hospitals and remote village clinics. The project highlights the importance of sustainability as a first-class goal for systems research.**

Since the end of World War II, substantial mechanisms for providing multilateral aid to underdeveloped nations have emerged including the World Bank and many regional and national development programs. A primary concern for projects supported by aid programs is sustainability. Although there is a growing belief that information and communication technologies (ICT) can significantly impact development,<sup>1</sup> in practice creating sustainable ICT projects is extremely difficult. A typical project consists of a pilot stage that aims to demonstrate the basic goals, followed by a deployment stage that aims for both scale and sustainability. Few projects even get to the latter stage, and many that do eventually fail.

The reasons for ICT failures vary, but at the core is an underappreciation of the many obstacles that limit the transition from a successful pilot to a truly sustainable system. In addition to financial constraints, there are operational issues such as power-related equipment failures, difficulties in local system maintenance, and an ongoing need for trained local staff, who often move on to better jobs. Further, researchers tend to focus on a deployment's sexier aspects, such as a novel technology, performance improvements, or a highly publicized pilot. But real impact requires a sustained presence in both deploying and maintaining ICT technologies, and

this implies that every aspect of the system must be designed to be financially and operationally sustainable within the local context.

## PRINCIPLES OF SUCCESS

Our experiences as members of the Technology and Infrastructure for Emerging Regions (TIER) research group at the University of California, Berkeley (<http://tier.cs.berkeley.edu/wiki/Home>), indicate that any ICT project must exhibit three important principles to be sustainable: optimization of an existing system, financial self-sufficiency, and operational self-sufficiency.

### Optimization of an existing system

Development projects are not deployed in a vacuum; there is almost always an existing approach for a given task, even if it is a poor one. We found it extremely valuable to view our work as optimizing the existing system rather than deploying something from scratch. There are three key reasons for this:

- the community understands the needs and motivation for the existing system, so variations on it require less communication and education;
- achieving community buy-in, which is fundamental to sustainability, is easier; and
- it is rarely clear to outsiders why the current system



is the way it is, which means larger changes might fail due to unforeseen circumstances.

For example, users in Ghana rejected a new medical-record system in part because it brought in too much transparency; lab workers secretly depended on performing lab tests on the side for extra income.<sup>2</sup> The optimization principle promotes a kind of humility that admits we cannot fully understand the existing problem space, and thereby prevents “blank slate” approaches that are unlikely to succeed in practice.

### Financial self-sufficiency

Any deployment aiming for sustainability must be cash-flow positive. Projects that fail to recover at least operating costs tend to fail overall, especially when outside funding diminishes at the end of a pilot. The goal of positive monthly cash flow is easier to achieve than profitability, which implies enough income to recover costs for capital investments. Although profitability is more desirable, maintaining positive cash flow is hard enough in rural areas, making it a reasonable target for development projects. Under this view, aid can be used for start-up costs but not for ongoing operations.

However, we can frame certain capital expenditures such as PCs as a monthly cost in addition to the operating costs, by dividing the up-front cost into equal monthly fees over the item’s expected lifetime. This calculation ignores the interest that money could have earned, which we estimate to be 8 percent. For example, a PC that costs \$600 and lasts five years translates to \$10 per month without interest and  $\$600 \times (1.08)^{5/60} = \$14.70$  per month with interest. But in developing regions, where old components are reused, the PC would not have zero value at the end of five years. Factoring in a salvage value of 20 percent reduces the monthly fee proportionally to \$11.76 per month.

### Operational self-sufficiency

A major operational issue is ongoing system maintenance and support. This includes power, hardware, and software as well as expansion and new installations. In practice, local groups do not start out with the ability to handle these tasks well. Also, some locations’ remoteness makes it hard for experts to help. Thus, three basic design goals for operational sustainability are

- increased component robustness,
- easy-to-use management tools for local staff, and
- tools for remote management by experts.

We have developed various techniques for these goals. As a complementary approach, we have also added a

middle tier, between local staff and remote management, which uses local IT vendors as a second level of support. We train both local staff and local vendors, but the latter have more IT background and can thus handle more issues and ultimately reduce the need for remote management. The vendors charge for their services, and we include these costs in our financial sustainability accounting.

### Example

Grameen Telecom’s Village Phone (VP) initiative, which provides rural phone access to more than 50,000 out of 68,000 villages in Bangladesh, illustrates all three principles.<sup>3</sup> VP is an example of optimization, as it simply extends the existing telephone system; it is financially sustainable, with profit on a per-minute basis; and it is operationally sustainable because operators learn how to maintain the phone and batteries and save money to replace hardware as needed.

**We found it extremely valuable to view our work as optimizing the existing system.**

### ARAVIND EYE CARE SYSTEM

In recent years, TIER has helped the Aravind Eye Care System ([www.aravind.org](http://www.aravind.org)) to deploy a rural wireless telemedicine system in southern India. Aravind comprises five eye hospitals in Madurai, Theni, Tirunelveli, Coimbatore, and Pondicherry in the state of Tamil Nadu. By volume, it is the largest eye-care provider in the world. In 2006-2007 alone, Aravind saw 2.3 million patients and performed 270,000 surgeries, most of which were for cataracts.<sup>4</sup>

Aravind’s stated mission is to eradicate needless blindness. The most common causes of preventable blindness are refractive errors and cataracts, which can be treated with prescription glasses and cataract surgery, respectively. With 15 million people needlessly blind, India has the largest share of preventable blindness globally.<sup>5</sup> About 70 percent of India’s population is rural, where the risks are higher but access to eye care is lowest.<sup>6</sup>

The primary limitation to eye care is a severe shortage of trained doctors and nurses in rural areas. India has only 10,000 ophthalmologists serving a population of 1 billion, with 90 percent of doctors based in urban areas.<sup>7</sup> Rural patients must typically travel long distances to clinics or hospitals. Travel expenses, even if not large in absolute terms, can be significant fractions of rural patients’ incomes, and as a result many are unable or simply decline to get treatment. We interviewed one patient who was in need of cataract surgery but had gone seven years without treatment until a local center opened; he was effectively blind that entire time.

### ‘Eye camps’

Aravind’s first strategy to address needless blindness was to conduct periodic “eye camps” in rural areas. A



Figure 1. Aravind vision centers, like this one in Periyakulam, are typically rooms rented from a rural family's home.

team of Aravind doctors would set out to a remote area and conduct comprehensive eye exams for an entire day. Although clearly an improvement, this approach reached at best only 7 percent of the target population.<sup>8</sup> The primary reason for low turnout was the difficulty of raising awareness among the community. The camps' transient nature also led to delayed treatments, low follow-up rates, and a limited ability to prevent problems. Finally, it did not really address the underlying issue of doctor scarcity.

**Vision centers**

To increase utilization of doctors, Aravind has adopted the *vision center* model, in which doctors remain at the hospital but interact with rural patients

over a communication network. A VC, shown in Figure 1, is typically a room Aravind rents from a rural family's home in the village. It is equipped with some basic ophthalmic equipment and a PC with a webcam. As Figure 2 shows, the center is staffed by two people: a technician who operates the ophthalmic equipment and PC, and a counselor who follows up with patients based on the diagnosis.

Center staff generally do not have a degree or a broad technical skill set; Aravind trains them specifically for their duties. At the VC, the technician performs some basic tests for refractive errors and cataracts. The counselor presents the results to the doctor at the base hospital via a videoconference, after which the patient interacts with the doctor. The counselor then follows up on the doctor's advice—for example, by handing out prescriptions, filling out referral forms, or creating glasses. If advised by the doctor, the counselor refers the patient to the base hospital for further examinations or treatments such as cataract surgery.

The cost to the patient for a VC consultation visit is 25 cents. Cataract surgery, if required, can cost up to \$75 at the hospital—surgery of comparable quality in the US costs about \$2,000. However, about two-thirds of patients cannot pay and receive surgery for free, which is also true for non-VC patients; these procedures are subsidized by paying patients.

**Wireless links**

In 2005, Aravind's eye hospital at Theni created three VCs based on corDECT wireless local-loop technology,<sup>9</sup> supplied by a local carrier focusing on rural connectivity. Each site, including the base hospital, had a total bandwidth of 36.5 kilobits per second. Not surprisingly, the video quality was insufficient, although the audio had some value. Going through a carrier limited Aravind's ability to start centers in areas with dire need. Despite being ready with clinical equipment and



Figure 2. Aravind VCs are staffed by two people: (a) a technician who operates the ophthalmic equipment and PC, and (b) a counselor who follows up with patients based on the diagnosis provided by a doctor at the base hospital.



personnel, it could not start VCs in two locations as the carrier did not consider those areas profitable enough to deploy a base station.

The same year, we established our own long-distance Wi-Fi link as an alternative connection for the Ambasamudram VC in Theni, mostly for operational experience. By early 2006, satisfied with the high performance of 5-6 Mbps per link and the operational freedom of an unlicensed spectrum, Aravind-Theni phased out the corDECT links, converting the existing three VCs and completing two others. Figure 3 shows the Aravind-Theni network's current status.

### Optimizing the system

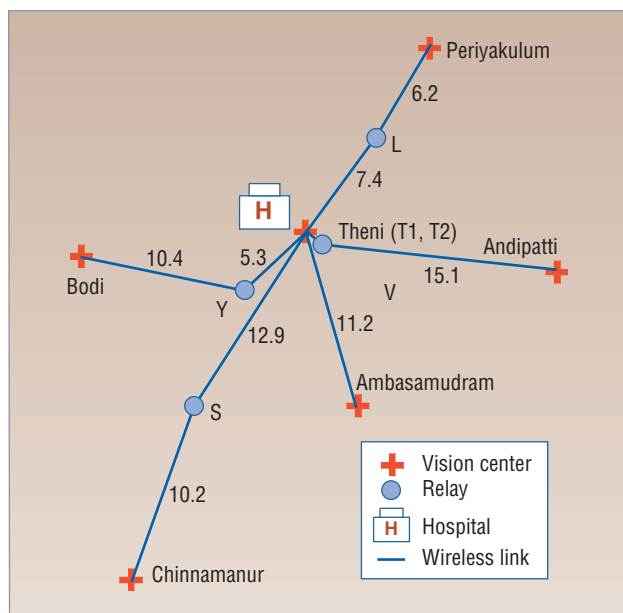
Our work at Aravind reflects the optimization principle: The doctors are great; we only needed to improve rural access to them.

Figure 4 shows the growth in the number of Aravind patients and VCs from January 2006 to December 2007. Ambasamudram, Andipatti, and Bodi were the existing centers that migrated to our high-bandwidth links, while the rest are new VCs enabled by our solution. Overall, the network enabled 51,205 remote eye examinations during that two-year period.

From May 2007, when all nine VCs were up, until December 2007, the system served an average of 3,632 patients per month. About 75 percent of patients visiting all VCs were new, while the remaining 25 percent came in for follow-ups. According to Aravind, this implies

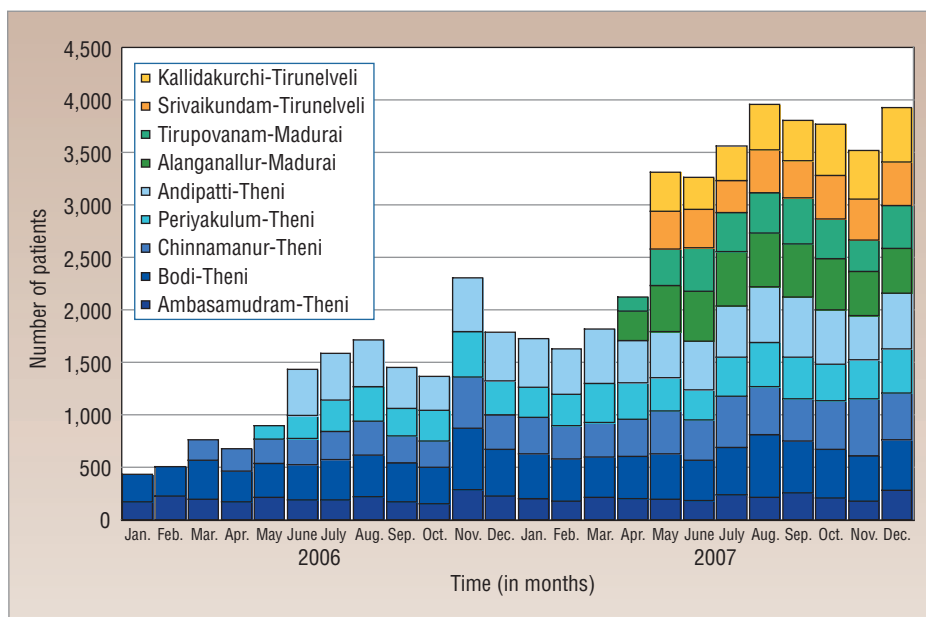
- more extensive eye care than the eye-camp approach because new patients get treatment every month, and
- higher-quality eye care with increased patient follow-ups.

Overall, 9,835 patients were diagnosed with severe cataract or refractive errors and needed significant vision improvement. Of these, 90 percent (8,814) got their sight back through prescription glasses or cataract surgeries as advised by the doctor during the videoconference. None of these villages have any ophthalmologists, except for Bodi, which has one doctor in private practice who visits once a week. We know from interviews that these patients generally would not have received treatment if not for the VCs.



**Figure 3. Aravind-Theni telemedicine network.** The base hospital at Theni is connected to five vision centers: Ambasamudram, Andipatti, Bodi, Chinnamanur, and Periyakulum. Theni has two wireless routers (T1, T2); the other routers are relays. All link distances are in kilometers.

These patients are also likely to return to income generation, the first step out of poverty. A recent study revealed that 96 percent of Aravind patients who get cataracts stop working.<sup>10</sup> Among those who lost their jobs, about 85 percent of men and 58 percent of women who get surgery return to wage-earning activities within a week.



**Figure 4. Patient throughput at Aravind vision centers.** The graph shows patients per month per center for two years. TIER set up the links to all Theni VCs; Aravind set up subsequent links to the new centers in Madurai and Tirunelveli.



## NETWORK DESIGN

Deploying high-performance networks economically in rural areas is fundamentally difficult for three main reasons.

First, Internet access of any kind is extremely limited, so there is not likely to be connectivity nearby except via satellite, which is very expensive.

Second, the typical rural landscape features low population density, with users spread across different villages separated by many kilometers. Therefore, deploying new network infrastructure is costly, with the cost amortized among only a small pool of users. This has led most broadband carriers to avoid these areas.

Third, users in rural regions have low purchasing power. Nongovernmental organizations (NGOs) deploying connectivity solutions are often financially limited, with finances proportional to potential customers' low income. Thus, the operating cost of network access is a limiting factor. Although aid often covers the initial costs of network access, the operating costs typically are not, or not for long.

### Wi-Fi-based solution

In our work with Aravind, we chose to deploy a novel Wi-Fi-based long-distance network as a low-cost connectivity solution.

**Advantages.** Wi-Fi reduces both deployment and operational costs in many ways. First, it does not incur frequency licensing fees. In addition, unlicensed frequency gives the hospital the operational freedom to put up links whenever and wherever needed, improving sustainability. Previously, Aravind could not add clinics in areas carriers considered unprofitable. Moreover, Wi-Fi equipment is low cost and low power.

One tradeoff we make is to provide islands of coverage using point-to-point links rather than blanket coverage for the whole region. This approach is clearly cheaper and seems to be sufficient, as the VCs are spread out and there is no need to support mobility. The blanket model also eliminates the ability to reuse spectrum, which reduces the effective local bandwidth.

Our solution comprises inexpensive, off-the-shelf hardware components. We use commodity Wi-Fi network cards, high-gain directional antennas and cheap, low-power single-board computers as Linux routers. The cost of a long-distance Wi-Fi link, excluding tower costs, is less than \$800.

**Limitations.** Wi-Fi is designed to work in high-density, short-range zones such as hotspots. Standard Wi-Fi makes little sense in point-to-point long-distance settings in rural areas with sparsely distributed populations. There are two main reasons for poor performance: protocol- and channel-induced losses.<sup>11,12</sup>

IEEE 802.11 relies on a simple stop-and-wait link-recovery mechanism. As the link distance increases,

the propagation delay increases as well, and the sender waits for a longer time for the ACK to return. This decreases channel utilization. Beyond the ACK-time-out limit, the sender may retransmit unnecessarily and waste bandwidth.

The protocol uses carrier sense multiple access with collision avoidance (CSMA/CA) technology, in which all stations listen to the medium before transmitting and send only when the channel is idle. Over long links, the state of the medium at the sender does not reflect the receiver's state. Thus, in the presence of bidirectional traffic, the sender and receiver cannot correctly assess the channel, leading to collisions. With increasing link distance, the probability that two end-hosts begin transmission within a window defined by correspondingly increasing propagation delay also increases, causing more collisions.

Another problem is that point-to-point links emerging from multiple radios colocated at the same wireless router interfere with one another when the directional antennas' high-energy side lobes overlap.

Along with these protocol shortcomings are channel-induced losses due mainly to external Wi-Fi interference. These losses can be highly variable and asymmetric. In addition to constant residual losses between 1 and 10 percent, there are bursty losses of varying magnitude and duration.

### WiLDNet

To address these issues, we have designed WiLDNet,<sup>11</sup> a time division multiple access (TDMA)-based media access control (MAC) protocol that operates on standard low-cost hardware. We improve link recovery through bulk acknowledgments instead of the standard stop-and-wait mechanism. We synchronize transmissions at colocated radios to mitigate the interlink interference problem, and combat residual losses by using forward error correction (FEC).

Our design provides a two-to-five-times increase in Transmission Control Protocol/User Datagram Protocol performance compared to the conventional MAC. We have used WiLDNet to achieve a bidirectional throughput of 6 Mbps over a single-hop 382-km link in Venezuela; to the best of our knowledge, this is currently the longest single-hop Wi-Fi link without any custom antenna design or active amplification.<sup>13</sup>

### DEPLOYMENT CHALLENGES

Even with appropriate network design, sustaining ICT deployments in underdeveloped regions is challenging for several operational reasons.

**Limited local expertise.** Most grassroots wireless network deployments are run by small groups with limited IT experience, which can lead to misconfiguration, limited diagnostic capabilities, and inadvertent equipment misuse.



**Hardware and software failures.** Poor-quality power causes various problems to low-cost equipment. Frequent power surges can damage the routers, Ethernet hubs, or power over Ethernet (PoE) injectors, while frequent power cycling can corrupt the software on the router's CompactFlash (CF) storage.

**Lack of connectivity.** Many problems can be investigated and solved by remotely logging in to the router. However, this assumes the network is working. When links failed, a frequent occurrence, we had to physically go on-site.

**Lack of physical access.** Traveling all day or night to get on-site is not uncommon. Worse, towers and relay points tend to be in uninhabited locations and might require hiking in. During many of our early deployments, we were forced to personally go into these remote places, and it took us days and considerable cost to deploy or repair equipment.

## POWER ISSUES

Any ICT network deployment must address power issues. Although power outages in rural India are well-known, we were surprised by the degree of power-quality problems. We measured voltage variations ranging from 70 to 350 V, with frequent spikes of 500 V, and some even reaching 1,000 V.<sup>13</sup> We lost quite a bit of equipment from these power fluctuations, as normal voltage in India is 220-240 V.

In general, we address availability via batteries or an uninterruptible power supply (UPS), both of which add cost. However, standard UPS systems are only the "standby" type; poor-quality power flows through untouched when available from the grid. Stable power is delivered only during an outage (when it comes from the batteries).

Table 1 shows the costs of the basic components of power management. Monthly cost is the capital cost divided by the lifetime; the last column includes interest on borrowed capital that is typically used to pay up-front costs.

### Cost of grid power

The real cost of rural grid power is overcoming the availability and quality issues. In India, a kilowatt-hour (kWh) of grid electricity costs about 5 cents on average.<sup>14</sup> A continuous 7-W draw for a wireless router requires about 5 kWh per month; thus, grid power costs about 25 cents per month.

The maximum outage length drives battery cost; assuming a 7-W draw for 10 hours a day, the requirement is 70

**Table 1. Costs of basic power-related equipment, in dollars.**

Power component	Capital cost	Lifetime	Monthly cost	Monthly cost with 8 percent interest
7-W rural grid power			0.25	
18-Ah battery	45	12	3.75	4.05
Small UPS without batteries	80	24	3.33	3.89
UPS with 18-Ah battery	260	24	5.21	6.08
Power controller	70	60	1.17	1.71
50-W solar panel	275	120	2.29	4.95
Controller + solar panel + 18-Ah battery	390	60/120/36	4.71	8.24

watt-hours per day. Using 12-V batteries, this requires about 6 amp-hours per day of outage, or about 18 Ah to tolerate three days without power. Battery life is quite variable and depends significantly on the frequency and depth of discharge. However, in practice, due to battery abuse from incorrect charging, the actual lifetime of batteries used directly is less than one year, and it is about two years for batteries used within a UPS.

Although raw electricity only costs 25 cents, a UPS with a sufficient battery for three days of outage costs \$6.08 per month, or about 24 times more. This represents the real cost of grid power. Some ICT projects incorrectly view the cost of electricity as zero because it is relatively common to steal electricity in rural India,<sup>15</sup> but such projects, unlike lighting or heating, generally need clean power.

### Power controller

To improve network uptime and increase component lifetimes, we have developed a low-cost power controller<sup>13</sup> that provides stable 18-V current to the wireless routers by combining solar and battery input. Its \$70 price is novel for the combination of features it offers such as peak power tracking, low-voltage disconnect, trickle charging, and PoE. Current chargers cost around \$300 and do not include all these features, which provide 15 percent more efficient power draw from the panel, can increase the battery lifetime by almost 200 percent, and enable remote management.

### Cost of solar power

As Table 1 shows, the full cost of solar power remains higher than grid power. We assume a 50-W panel with a lifetime of 10 years and a 5-year lifetime for the controller and the same 18-Ah battery. Due to the panels' cost, the overall cost is \$8.24 or about 35 percent higher than grid power. This \$2 difference might be acceptable given the better power, but the difference scales with total power. Absent financial gain, the main value is the ability to obtain reliable power essentially anywhere, which is particularly useful for remote relay points.



**Table 2. Operational expenses, in dollars, for Aravind-Theni vision centers.**

	Ambasamudram		Bodi		Chinnamanur		Periyakulum		Andipatti	
	2006	2007	2006	2007	2006	2007	2006	2007	2006	2007
Local income	3,313	3,590	5,526	5,984	3,855	6,381	3,837	5,311	7,110	8,392
Referral income	1,534	1,890	1,928	2,550	977	2,025	1,018	1,845	2,145	2,865
Expenses	4,676	5,704	6,799	7,377	6,071	6,354	4,530	6,973	8,032	8,451
Cash flow	171	-224	655	1,157	-1,239	2,052	325	183	1,223	2,806

## SYSTEM MANAGEMENT

In addition to power issues, any ICT deployment must also address overall systems management issues of monitoring, remote management, automated local recovery, and ongoing training. TIER's primary goal is to increase system availability, decrease operational expenditure, and ultimately transfer operational responsibility to local teams.

### Three-tiered support system

To leverage the range of skills and availability of potential support personnel, we have created a three-tiered support system.

Tier 1 consists of *local staff* responsible for basic management and maintenance. Tier 2 includes *local network integrators*—local vendors trained in installation, configuration, and debugging of networking components. Tier 3 consists of the *remote management team*, comprising of highly skilled professionals familiar with all the hardware and software. At Aravind, TIER has played the role of the remote management team.

### Components

We have built a monitoring system, alternative backchannels for remote management during link failures, and automatic-recovery mechanisms that together have improved operational sustainability.<sup>13</sup>

For monitoring, we implemented a “phone home” system that pushes data logs to our US-based servers every three hours. The system collects router-level, link-level, and end-to-end measurements. In addition, it enables interactive remote management by opening reverse Secure Shell (SSH) tunnels to the Aravind routers.

Link failures leave the network partitioned. In many cases, alternate access to the routers can fix simple problems such as misconfiguration and can prevent trips to router locations. We have created hop-by-hop backchannels independent of the network configuration using link-local IP addressing.

At the base level, we provide a suite of simple mechanisms that enable the system to automatically recover from basic and localized errors such as a hung OS, disk errors, or stalled boots owing to low voltages. An important feature of automatic recovery is a watchdog-

processor-based triggered reboot of the system upon detecting a software malfunction.

Our three-tiered support system uses all these components. Tier 1 personnel have a simple monitoring interface that presents the liveness of all network links and nodes and corresponding usage statistics. Tier 2 and Tier 3 personnel have access to detailed performance information including end-to-end throughput characteristics, signal-strength measurements, power levels (voltages and currents), interference levels, and the number of reboots and their possible causes. All personnel also have backchannel access wherever available for remote management.

## FINANCIAL SUSTAINABILITY RESULTS

Table 2 shows the earnings of all five VCs in Theni. Local income includes patient payments for consultations, medicines, blood-sugar tests, spectacles, and return of excess inventory. Referral income includes 20 percent of the cost of cataract surgeries at the base hospital for patients referred from the center. Expenses include various office expenses and staff salaries, including payments to the doctor at the base hospital for consultations, but do not include network costs, which we discuss later.

To achieve positive cash flow, VCs must get some of the hospital income for referrals. This is reasonable but implies that the relatively few expensive procedures are part of the centers' sustainability. If referral income is included, four out of five VCs were cash-flow positive in 2006 and 2007; without it none were positive in 2006, although Chinnamanur managed to be positive, and Andipatti was borderline in 2007. These results have encouraged Aravind to set the bar higher and not rely on referral income for sustainability in the future as they believe it is now achievable.

We now include the network costs. As an example, consider the Bodi VC, which made about \$650 and \$1,150 in 2006 and 2007, respectively. The local vendor that maintains the links charges about \$18 per link per year for maintenance, excluding hardware replacement costs. Bodi uses two hops with a relay for an annual cost of \$37.70. The capital cost for the two hops was about \$1,800, including part of a shared tower. Assuming a five-year lifetime for everything except the tower, a 20 percent



salvage value, and 8 percent interest, this translates to \$32 per month, or about \$385 per year. Thus, Bodi remained cash-flow positive in 2006 and 2007 after including both the capital and operating cost of the network.

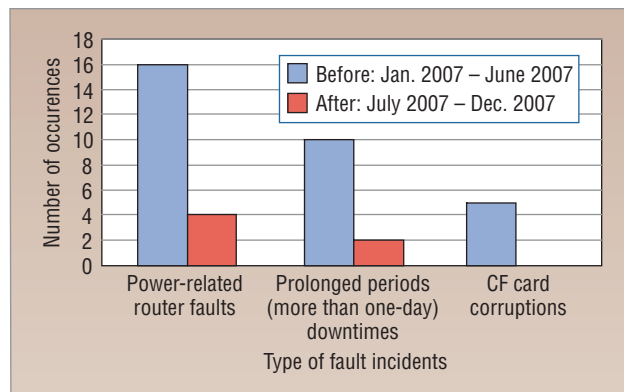
The total capital cost for a VC is \$11,000. This includes ophthalmic equipment, furniture, the PC, the network, and outfitting the building. If we assume half the value remains after 10 years, the monthly cost with interest is about \$100 per month, or \$1,200 per year. This implies overall net losses of roughly \$500 in 2006 and \$40 in 2007 for Bodi. The center is not truly profitable but it is close. For example, assuming 65 percent salvage value after 10 years and a 4 percent interest rate leads to breaking even.

Thus once created, VCs appear to be sustainable. However, the hospital is not likely to recoup its initial capital outlay, nor is it earning enough from these centers to pay for the next ones. Nevertheless, Aravind views the centers as sustainable enough and is expanding the system to 50 centers. Because the VCs are on the edge of viability, we expect Aravind to maintain its emphasis on controlling operating costs and extending equipment lifetimes, and on using aid money for new centers, which matches the mode of other development groups. On the positive side, the number of patients per month continues to rise, which increases financial sustainability.

### OPERATIONAL SUSTAINABILITY RESULTS

Our solutions to power issues and for overall system management have improved the Aravind-Theni network's operational sustainability. As Figure 5 shows, since improving power quality we have reduced the number of power-related router downtimes. We have also managed to limit downtime periods to under one day; this is primarily due to quicker diagnosis enabled by backchannels. Through a combination of quality power and software safety features such as a write-protected OS, we have eliminated CF card corruptions.

Over time, operational responsibility for the network has migrated from us (Tier 3) to local staff (Tier 1), and we consider this a big success. As Figure 6 indicates, we were solely responsible for the entire network. Through training and development of management tools, local staff has learned to maintain and manage the network on their own, while a local vendor (Tier 2) handles tower construction, antenna alignment, and other installation issues. In the past year we have not installed any links ourselves even though Aravind has established four additional VCs. Our role is reduced to supplying equipment for new wireless installations,

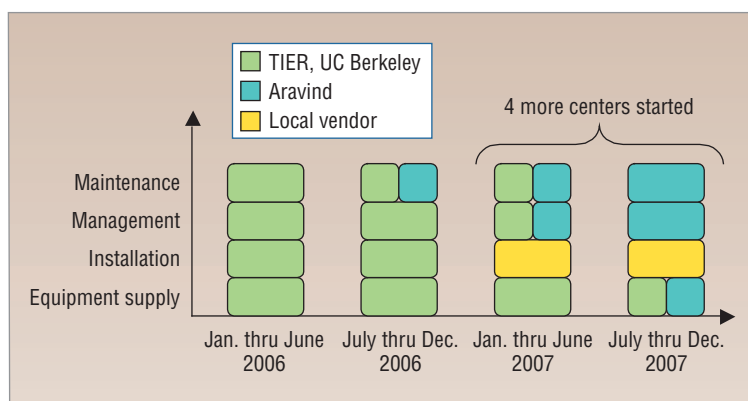


**Figure 5. Reduction in Aravind-Theni network faults. Introducing stable power, reliable recovery, backchannels, and safe software improved the network's operational sustainability.**

which Aravind now also pays for as the centers have demonstrated they can recover network costs.

**A**ravind's vision centers have been a profound success for both the hospitals and the patients involved. In two years of operation, the system has enabled over 50,000 remote patient-doctor consultations; the current rate is 3,600 per month. For more than 8,000 patients, the VCs led to restoration of their vision and new economic opportunities.

To continue this success, the centers must be both financially and operationally sustainable. We worked to achieve this by reducing deployment and operating costs, increasing deployment flexibility, and developing a long-term support plan that creates a local ecosystem and transitions to local staff and vendors over time. We also developed a range of new tools for system management and to address power issues. Our choices had a large practical impact on making



**Figure 6. Migration of operational responsibility for the Aravind-Theni network. Initially responsible for all network aspects, TIER now only supplies wireless equipment for new VCs. Local staff have learned to maintain and manage the network on their own, while a local vendor handles installation issues.**





the VCs cash-flow positive, which previous attempts could not do.

Aravind is now scaling the Theni network to 50 centers based on the success of this effort. The full network is expected to handle 500,000 exams per year and provide eye care to a rural population of 2.5 million within the next three years. Currently, the VCs break even, but are not quite truly profitable and would have a hard time paying back start-up loans or generating capital for future centers. Although Aravind feels that the network is “close enough” to sustainability, walking this tightrope highlights the ongoing challenge of sustainability and why it must be a first-class goal for systems research. ■

---

### Acknowledgments

We thank the Aravind Eye Care System for their partnership. We also thank the editors, Bernardine Dias and Kentaro Toyama, for their help, and Yahel Ben-David, Alan Mainwaring, Lakshminarayan Subramanian, and Manuel Ramos for their opinions, data, and assistance on a previous draft. This work was supported by National Science Foundation Grant No. 0326583 and Intel Research.

---

### References

1. E. Brewer et al., “The Case for Technology in Developing Regions,” *Computer*, June 2005, pp. 25-38.
2. Interview with Melissa Ho, PhD candidate, UC Berkeley School of Information; [www.ischool.berkeley.edu/people/students/melissaho](http://www.ischool.berkeley.edu/people/students/melissaho).
3. N. Cohen, “What Works: Grameen Telecom’s Village Phones,” Digital Dividend study, June 2001, World Resources Institute; [www.digitaldividend.org/pdf/grameen.pdf](http://www.digitaldividend.org/pdf/grameen.pdf).
4. Aravind Eye Care System, *Annual Activity Report: 2006-07*; [www.aravind.org/annualreport](http://www.aravind.org/annualreport).
5. L. Dadona, R. Dadona, and R.K. John, “Estimation of Blindness in India from 2000 through 2020: Implications for the Blindness Control Policy,” *National Medical J. India*, Nov./Dec. 2001, pp. 327-334.
6. P.K. Nirmalan et al., “Utilisation of Eye Care Services in Rural South India: The Aravind Comprehensive Eye Survey,” *British J. Ophthalmology*, Oct. 2004, pp. 1237-1241.
7. G.V. Murthy et al., “Human Resources and Infrastructure for Eye Care in India: Current Status,” *National Medical J. India*, May/June 2004, pp. 128-134.
8. A.E. Fletcher et al., “Low Uptake of Eye Services in Rural India: A Challenge for Programs of Blindness Prevention,” *Archives of Ophthalmology*, Oct. 1999, pp. 1393-1399.
9. A. Balaji, T.A. Gonsalves, and R. Kumar, “The corDECT Wireless Local Loop,” *Proc. 1999 IEEE Int’l Conf. Personal Wireless Communication (PWC 99)*, IEEE Press, 1999, pp. 359-363.
10. K. Gooding, “Poverty and Blindness: A Survey of the Literature,” briefing paper, Feb. 2006, Sight Savers Int’l; [www.icch.org.uk/files/poverty\\_and\\_blindness\\_02.06.pdf](http://www.icch.org.uk/files/poverty_and_blindness_02.06.pdf).
11. R. Patra et al., “WiLDNet: Design and Implementation of High-Performance Wi-Fi-Based Long-Distance Networks,” *Proc. 4th Usenix Symp. Networked Systems Design and Implementation (NSDI 08)*, Usenix Assoc., 2007, pp. 87-100.
12. A. Sheth et al., “Packet Loss Characterization in Wi-Fi-Based Long-Distance Networks,” *Proc. 26th IEEE Int’l Conf. Computer Communications (Infocom 07)*, IEEE Press, 2007, pp. 312-320.
13. S. Surana et al., “Beyond Pilots: Keeping Rural Wireless Networks Alive,” *Proc. 5th Usenix Symp. Networked Systems Design and Implementation (NSDI 08)*, Usenix Assoc., 2008, pp. 119-132.
14. P. Audient and F. Verneyre, *Electricity in India: Providing Power for the Millions*, Int’l Energy Agency, 2002.
15. M. Gregory, “India Struggles with Power Theft,” *BBC News*, 15 Mar. 2006; <http://news.bbc.co.uk/2/hi/business/4802248.stm>.

*Sonesh Surana is a PhD candidate in the Electrical Engineering and Computer Sciences (EECS) Department at the University of California, Berkeley. His research interests include low-cost wireless infrastructure, mobile phone applications, and sustainable technology deployments in developing regions. Surana received an MS in computer science from UC Berkeley. Contact him at [sonesh@cs.berkeley.edu](mailto:sonesh@cs.berkeley.edu).*

*Rabin Patra is a PhD candidate in the EECS Department at UC Berkeley. His research interests are in low-cost wireless networking and other ICTs for developing regions. Patra received a B.Tech in computer science and engineering from the Indian Institute of Technology, Kharagpur. Contact him at [rkpatra@cs.berkeley.edu](mailto:rkpatra@cs.berkeley.edu).*

*Sergiu Nedeveschi is a PhD candidate in the EECS Department at UC Berkeley. His research focuses on wireless and energy-efficient networks, and the economic and social aspects of computer networking in developing regions. Nedeveschi received an MS in computer science from UC Berkeley. Contact him at [sergiu@cs.berkeley.edu](mailto:sergiu@cs.berkeley.edu).*

*Eric Brewer is a professor of computer science in the EECS Department at UC Berkeley and the director of Intel Research Berkeley. His primary research interests include Internet systems, developing regions, sensor networks, and programming languages. Brewer received a PhD in electrical engineering and computer science from the Massachusetts Institute of Technology. He is a Global Leader for Tomorrow for the World Economic Forum, a member of the National Academy of Engineering, and an ACM Fellow. Contact him at [brewer@cs.berkeley.edu](mailto:brewer@cs.berkeley.edu).*