

The Challenges of Technology Research for Developing Regions

Our projects in several developing regions ran into many technical, cultural, and environmental challenges. We hope to help others, especially researchers from outside these regions, avoid our mistakes.

Eric Brewer, Michael Demmer,
Melissa Ho, R.J. Honicky,
Joyojeet Pal, Madelaine Plauché,
and Sonesh Surana
University of California, Berkeley

The work we've done in the world's developing regions¹ has been immensely rewarding and helped make us all strong advocates for technology. The first reward is simply the experience of these countries, their surprising happiness, and the technology insights that come from being there. The deeper reward is of course actually helping people by creating new options through technology. For example, a video-conferencing solution on top of our novel low-cost long-distance networking has directly led to improved eye care for about 30 patients a day in our test site in rural India. Although we've worked in eight countries, most of our examples come from four projects:

- In the telemedicine project just mentioned, with the Aravind Eye Hospital in Tamil Nadu, India, our goal is to diagnose patients that are too far from the hospital to visit in person (www.aravind.org).
- In the village kiosks of the M.S. Swaminathan Research Foundation (MSSRF), also in southern India, our goal is to use speech recognition for semiliterate villagers (www.mssrf.org).
- In a rural networking project in Ghana, we're looking at low-cost wireless backhauling of IP traffic (that is, carrying aggregate Internet traffic for a group, such as an ISP).
- In an Asia Foundation project in rural Cam-

bodia, we're focusing on improved email and content distribution applications over very poor networks (www.asiafoundation.org).

However, we encountered a wide range of technical, environmental, and cultural challenges that are outside the scope of typical computer science research.

In this article, we share some of our experiences with the hope of increasing understanding of these issues. We also hope to help others, particularly researchers from outside these regions, to avoid our mistakes. We document real challenges (mostly our own), try to generalize them, and suggest steps that might at least mitigate the problems.

Technical challenges

Although the specific technology issues we encounter differ somewhat from one developing region to another, they are in many ways easier to solve (than environmental or cultural challenges). Our problems ranged from equipment and power failures to software and remote management.

Equipment failures

One of the first lessons we learned is simply that typical equipment specifications aren't realistic for developing regions. The most obvious issue is humidity, which can often lead to condensation inside the equipment, eventually shorting the circuitry. For pole-mounted routers, we designed a box with a slanted top, so that condensation flows down the slant and around the board rather than dripping on it.



Figure 1. A dusty UPS system in Cambodia.

Heat is the next obvious problem. In developing regions, equipment must operate in a wide range of temperatures because of the lack of widespread climate control. Although cooling the rooms containing computers is sometimes possible, doing so is expensive and requires power, and cooling systems are unlikely to be handled by a UPS (uninterruptible power source).

Dirt and dust are also an issue (see figure 1). In Pailin, Cambodia, we found that the staff of a rural center had to clean the inside of computers once a week using an electric blower. The amount of dust accumulated was incredible, owing to the center's location in an open-fronted office on a relatively busy dirt street. Without this intervention, the machines would become unreliable and stop functioning over time.

All of our outdoor equipment needed to be especially robust to water and ultraviolet exposure. In one example, we installed a makeshift plastic enclosure in Pondicherry, India, as a temporary measure, but exposure to the elements rendered the plastic so brittle that the equipment was ruined in less than six months. UV protection in particular is both critical and hard to get; for example, we had

trouble finding UV-rated CAT5 cable in several countries.

In general, we lose equipment on every trip. Although power tends to be the biggest short-term cause, equipment also breaks because of bumpy roads. One system stopped working immediately after a very bumpy auto rickshaw (three-wheel mini taxi) ride over an 11-kilometer dirt road. We can't be sure that the ride was at fault, but one of our laptops also stopped working on that ride.

Power problems

The obvious problem with power is outages, both long and short, and UPS backup is the usual solution. However, the power in some centers in India and Cambodia was so intermittent that the UPS would frequently turn on and off, ultimately reducing battery life. So, in at least one case, when the power finally did go out for a longer period of time (more than a few seconds or minutes), the UPS battery was already dead. For the wireless backhaul project in Ghana, we determined that battery backup for the grid power had to last for at least three days, implying a relatively expensive battery system. For these types of networks, power is the second most expensive ele-

ment, after the cost of the towers (if there are any). These costs are compounded further in multihop networks.

Even when the power is on, the quality is often insufficient for IT, which is more demanding than the typical usage for lighting and cooking. Using a power logger, we observed spikes of up to 1,000 volts on a 220-V grid, as well as long periods of 170 V (brownouts). Anecdotally, AC adapters such as those for cell phones end up having very short lives, occasionally even less than six months. Many UPS systems only help a little, as they pass through the grid power and its problems, switching to battery mode only when the power is completely out.

On the positive side, laptops have somewhat fewer problems because of their integrated batteries, and the more expensive UPSs tend to work somewhat better. Finally, most middle-class homes and all data centers use gas or diesel generators for backup power, so this is an option, albeit an expensive one.

Intentional power outages also caused problems. In one Tamil Nadu center, our proxy cache was in an error state every morning, requiring a manual reboot. After some investigation, we found that the operators shut off the entire center's power each evening—a fairly common practice in many places. Yet it was still not obvious why the cache ended up in a hung state rather than rebooting cleanly, as we had specialized watchdog hardware that should have automatically rebooted the system. The real problem was that a single master switch turned on the whole center at once, leading to high current demand and a subsequent voltage drop. Booting this particular hardware board with insufficient voltage caused it to hang without enabling the hardware watchdog. Our current solution is still to reboot manually every morning, but a simple time-delay circuit to offset the boot times would be a nicer solution.

Software infection and reinstallation

Many of the Windows PCs used at Aravind and MSSRF were poorly maintained and infected with spyware and viruses. This meant that several applications, including the new ones we tried to install, didn't work properly. Similarly, general Internet access in Cambodian centers meant that viral infection was basically inevitable, since no one uses firewalls or antivirus software and most users do not understand viruses or spyware. As we later found out, the standard procedure for dealing with these computers' problems is to fully reinstall and reconfigure Windows. This is due in part to a lack of technical know-how to implement more targeted solutions, but in many cases it's the only safe solution.

In some cases, the impact on our project came down to this: the monitoring software we had installed was erased within a week of our departure because we had failed to leave sufficient reinstallation instructions. In the case of Aravind, unlike in Cambodia, we left reinstallation CDs at each site, which seemed to help significantly—at least our software got reinstalled along with the rest of the system.

However, we also had to configure the PCs for specific network settings and routing-table entries. Such site-specific settings get lost during reinstalls, so they must be reconfigured. One lesson here is that we should treat the OS installation as soft-state, by which we mean it's temporary, unreliable, and easily restored to a safe version. In turn, this implies we must gather any data and complex configuration onto separate storage systems, such as a local server or USB key. Also, we did not have reinstallation issues with Linux systems because they generally were left alone and tended not to get infected.

Remote management

Remote management seems like an

ideal solution to aid research in distant places. Local contacts are often untrained to diagnose or support problems, and it's clearly not time- or cost-effective to repeatedly fly graduate students to remote locations. Unfortunately, remote management has been extremely difficult to achieve in practice.

Given a good connection, we can securely log in to remote machines via the Secure Shell protocol (SSH) and manage or upgrade them. Most of the trouble is

how to ensure a good connection. Network Address Translation and firewalls bar incoming connections, and for satellite or dialup-based connectivity, the link is simply down until an outgoing connection initiates connectivity.

The basic solution is to have the remote device "phone home": we initiate connectivity from the remote side and then use the opened connection and SSH to log back into the device. However, this approach often works for a while and then stops. Problems we've found include a lack of power at the scheduled time, hung nodes (despite a watchdog), and loss of the capability because of misconfiguration or reformatting of the entire machine.

Once "phone home" stops working, we tend to be in trouble, as our machines are on the far side of the planet without network access. We're considering "sneaker net" style solutions such as USB keys or CDs that we can ship to remote sites to auto-install a patch, but this isn't easy, especially when we have limited information about the actual problem.

Finally, beyond simply enabling con-

nectivity to our machines, we must design systems that can revert back to a safe state, including correct local configuration settings. A smarter variation of a safe-mode boot or a custom reinstallation CD with all known (good) settings for a specific host might suffice.

Environmental challenges

Logistical problems include transportation, customs and shipping, and local production. There are also height-

Once "phone home" stops working,
we tend to be in trouble, as our machines
are on the far side of the planet
without network access.

ened risks of personal safety in urban slums and natural disasters.

Transportation

Travel time often hinders research, especially in rural areas. Kwaku Boadu, a wireless-network CEO in Ghana, said that travel costs are a significant part of operating expenses and limit a company's reach. Average speeds of 20 km per hour are common, and we've had many all-day or all-night trips.

Worse, each trip to a remote location has a nontrivial risk associated with it. To minimize risk to his engineers, Boadu avoids hiring professional drivers, realizing that they'll feel obligated to drive even when severely tired. Similarly, while working in Tamil Nadu, India, two of us narrowly avoided a serious accident when our van's driver had a seizure. Luckily, he was able to stop the van (somewhat roughly) before blacking out and convulsing. The driver had no known history of seizures, so the best explanation we could find was that he (a poor man) had eaten spoiled food the previous day and nothing else the



Figure 2. This pickup truck was stuck on a bridge in Cambodia that heavy rainfall had damaged.

following morning. Although ill, he apparently felt that he could not afford to miss work.

Weather and other unpredictable factors often exacerbate transportation problems. In Cambodia, we had to cancel a trip to a remote, rural information center in Preah Vihear because of an impassable road (see figure 2). We had assumed the road was still good because our driver had traveled it just a few months earlier—but heavy rains intervened.

Transportation problems thus increase the need for reliable equipment and remote management, both for cost and safety reasons.

Customs and shipping

Whenever possible, we bring all needed equipment with us or arrange to buy it locally, but for various reasons, sometimes we must ship it. However, shipments can be delayed while clearing customs and afterward.

Officials who were confused by the labeling of our wireless-equipment shipment incorrectly thought that it might violate spectrum usage policy. For 10 days, our equipment sat idle at the airport while we clarified the situation. We now try to use only simple, innocuous labeling and also find that shipping to Intel India works better than to less-well-known receivers. To avoid delays while we're working in a developing region, we try to ship our equipment a month in advance.

In India, we have to pay customs' duty of about 20 percent. Although import taxes on computer equipment are generally coming down, they remain high in many countries. Clearly, these customs and tax issues imply that shipping isn't a good long-term strategy, but it does simplify first-time deployments, after which we can invest in local procurement.

Local purchase and manufacturing

In the long term, we intend to transfer the technologies that we develop to local firms. This helps with both cost and scalability as well as encouraging local economic growth.

With that in mind, we've begun to identify local manufacturers and vendors who can provide equipment and skills. For our backhaul link in Ghana, we hired a man who makes satellite dishes to build some high-gain Wi-Fi antennas. Although some of his tools are rudimentary by modern standards, he clearly understands how to adjust a dish's geometry to different sizes and what gain to expect from a particular-sized dish at a certain frequency. Satisfied with his craftsmanship, we ordered two custom dishes, tailored to the 2.4 GHz band.

Building the dish was fast and successful, but when it came time to mount it, he couldn't weld the mounting brackets because the power was out for several days. He purchased a generator, but it had insufficient power for the arc

welder. Worse, the next day, he came down with malaria. Despite his illness, with our departure date quickly approaching, he drove the finished dish across town to a place with power and welded the mounting bracket for us. On the last working day of our visit, we hoisted the dish onto the mast and configured the network.

It was worth the wait. After some configuration tweaking, the dish yielded approximately 31 decibels (close to the theoretical maximum of 33.4 dB for a dish that size and shape), making a solid connection supporting 7 Mbps on an 802.11b link over 56 kilometers with virtually zero packet loss at the link layer. We recently installed the second dish on a return trip. This example reflects all the issues: strong local knowledge and craftsmanship, limitations to productivity, and unexpected delays.

Developing regions vary widely in the quality of local manufacturing, from high-quality, low-cost items to poor-quality, (relatively) expensive items. For our second trip to India, we knew that Ethernet crimpers, jacks, and cables would be readily available in Chennai's huge high-tech bazaar. However, we discovered that some crimpers and jacks were so badly manufactured as to be basically unusable. For the RJ45 jacks, pushing the Ethernet wires into their slots in the right order was almost impossible. We are now more careful in making local purchases. Finally, local procurement takes considerable time because of communication and transportation problems and the need to visit many locations. Despite these issues, local purchases remain preferable to shipping equipment in from outside.

Urban danger

Although we've done most of our work in rural areas, some was in the urban slums of India and Brazil. Espe-

cially in Brazil, we found that research could be a perilous task. Our work in Rio de Janeiro's slums, known as favelas, revolved around using refurbished computers in subsidized public kiosks. Basically, no research can occur in the favelas if the local drug lords choose to prohibit it. Even entering or leaving a favela without a local escort is usually out of the question. In some favelas, hiring someone to do the surveys was difficult because few agencies have experience working in the poorer neighborhoods. Getting a good sample of computer users was difficult because door-to-door visits were too risky, so we had to make do with sampling respondents at the favelas' outskirts rather than in the labyrinthic interiors. In the two-week survey-taking period, we had one incident in which a surveyor walked into a drug-related shootout.

Taking photos was very difficult, and the "informal" authorities prohibited it. Pictures contribute to creating street maps, which are strategic during drug raids. We were repeatedly told that drug lords and local gangs monitor computer center activities. One center was always warned in advance of shootings to help keep the computer students out of danger.

As an aside, some photos we took of graffiti in a Sao Paulo slum became a tool for threats between gangs. After one of us posted these pictures on a personal Web site, they surfaced in Google queries for one of the gang's names. We found out later that Brazilians, probably gang members, had used the Web site's comment features to send coded messages and threats to each other. Somewhat surprisingly, the messages' authors took the trouble to mask their IP addresses.

Natural disasters

On one of our trips to India, we arrived shortly after the December 2004 Indian Ocean tsunami, which had nearly destroyed two of the villages where we were working. Despite advance arrange-

ments, when we arrived we were told that we couldn't go to the villages—so many news crews had been in the area that the villagers were resentful. Although we used our time doing a combination of relief work and unplanned research, we were completely unable to address our original research goals.

You might be tempted to ignore disasters as a research challenge, given that they're rare. However, life is generally more fragile, less prepared, and more

overcrowded in developing countries, all of which exacerbate crises. Just in the past two years, we've seen the tsunami, a devastating earthquake in Pakistan, and incredible flooding in Bangladesh. Less damaging but more frequent are local crises, including broken dams, school fires, and monsoons. These events are systemic problems that, although unpredictable, will likely affect field work at some point.

Cultural challenges

Cultural issues present the greatest challenges, and they vary widely by region. We've faced IT staff problems, tampering and theft, corruption, and illiteracy.

Staff and training

Overall, developing regions have severe shortages of IT personnel, who can often make more money in other countries or in their own country's wealthier cities (such as Bangalore).

When deploying our work into our partners' infrastructures, we must work closely with their IT support staff. How-

ever, these staff members are typically overburdened. At the Theni Aravind Hospital, the network administrator was also responsible for data entry of patient records, statistics for all patients visiting the clinics and the eye camps, and weekly reports. He was hard-pressed to find time to maintain the existing network, let alone to help us install new technologies.

We've also worked with staff people who were relatively inexperienced or lacked confidence. The hospital's tech-

Cultural issues present the greatest challenges, and they vary widely by region. We've faced IT staff problems, tampering and theft, corruption, and illiteracy.

anical staff could do basic Windows configuration, but they couldn't help with more complex tasks such as modifying routes on their main gateway. Furthermore, lack of Linux knowledge meant that the staff was essentially unable to follow even detailed instructions about SSH, routing tables, and software installation. Similarly, after we set up the network, a staff person would occasionally contact us to say that the link "wasn't working." However, in many cases the root cause was something simple, such as a loose Ethernet cable or a mistyped ping command.

Many organizations also face high IT staff turnover—as the staff become more experienced, they leave for more lucrative positions elsewhere. At the Madurai Aravind Hospital, our project was delayed when the experienced network administrator with whom we were working left for a more lucrative job in Chennai. Much of the project knowledge left with him, and now we must invest a substantial amount of time in training a new administrator. In fact, the Aravind hospitals face similar problems



Figure 3. A battery in a locked cage in Ghana. Heat dissipates better from the battery in a cage than in a lockbox.

with their doctors as well. No doctors go to the city of Theni permanently, despite incentives such as increased salaries. However, in Aravind's case, the management views this training and turnover as a positive contribution: they see themselves as creating more doctors for the region and contributing to many young doctors' careers.

Where Windows is prevalent, it's hard to find experienced people to help us maintain Linux-based equipment remotely, and it's hard for us to provide that training. To mitigate this problem, we've had to develop detailed training manuals and provide a Web-based user interface for viewing and modifying the configuration. These reduce the learning curve and also simplify remote management but add more delays to the research efforts.

Finally, training staff to troubleshoot is important, so that when they encounter a problem they can't solve, they at least can give us more information than "the link isn't working." We have the local admins accompany us during the initial network setup, and after training we leave behind spare parts for them to use.

Recently, we've had some success asking local admins to set up links themselves as we watch or stand by on the phone; this breeds confidence in the admins and seems to bring them up to speed more quickly. We're also starting to partner with faculty or students in neighboring universities to act as local contacts for networking issues.

Tampering and theft

Fabio Rosa's company installs solar power for people's homes off the grid in Encruzilhada, Brazil, renting the equipment at the same cost as they would normally spend on candles, batteries, and lamp oil.² Initially, his team had to return often to make repairs, usually because people became curious about the equipment left in their home and started tampering with it. The solution? He asked each family who their favorite saint is and put the corresponding figurine on the lid of the box containing the equipment, thus preventing them from disturbing the equipment and underscoring that they should leave it alone.

In Cambodia, we deployed packet-

level network-monitoring software to log data onto a USB drive; this lets us retrieve the data later, swap out the drives when full, and avoid filling local hard drives. We later found that the USB drive's value was too much of a temptation—several local center operators had appropriated some drives for personal use, and others were simply missing.

Perhaps a better solution would have been to make a deal with the center operators when we were installing the software, such as giving them several USB sticks up front in exchange for helping us collect the data. We needed such an agreement with the operators in addition to a broader agreement with the local partner, which we did have. Typically, the operators value the actual hardware whereas the researchers value the data it contains.

Batteries are also regularly stolen, as the cage in figure 3 implies. In Ghana, Arrow Networks found that one of its remote hilltop towers stopped working, only to find that the batteries were gone. Their solution: relocate a farmer onto the tower's land in exchange for guarding it.

Corruption

Although we know of corruption problems in many countries, both industrialized and developing, our primary direct experience occurred in Ghana. Working with Arrow Networks, we're building a Wi-Fi-based backhaul network between Accra (the capital) and Kumasi in the north, about 300 km in total.

Two of the towers that Arrow Networks uses on the route from Accra to Kumasi are owned by Ghana Broadcasting Company and used for FM radio broadcasting. Arrow Networks had an existing contract with GBC to arrange to turn off the transmitter when engineers needed to climb the tower, for a fee of US\$100 per visit.

When we arrived, however, GBC's

Figure 4. Collecting Tamil speech samples from a semiliterate speaker in India.

regional director in charge of that particular tower raised the fee to a million cedis (US\$660) per hour. Clearly this violated Arrow's contract, but he was obstinate, and contracts are difficult to enforce in Ghana.

As it turned out, a personal contact who happened to be a radio personality for GBC later contacted the director and explained that our project was intended to bring low-cost connectivity to rural Ghana. The director then switched his position, saying that we could use the tower for free.

Ghanaians have a strong sense of social responsibility and also want to help out their associates. We were lucky to be able to tap both: the personal contact and the social goal both probably contributed to the director's change of heart.

Illiteracy

In many developing regions, illiteracy rates above 50 percent are common,³ making computers and technology inaccessible without a literate mediator's help. Many researchers have suggested speech recognition as a key to universal access, but success stories are rare. In addition to the challenges of dialectal variation, multilingualism, cross-cultural barriers, and lack of linguistic resources, we found that illiteracy presents a significant obstacle to the traditional techniques of recruiting, recording, and user testing.

Recruiting. In Tamil Nadu, recording illiterate speakers saying digits (zero to 10) took approximately six times as long as it did for literate speakers. The discrepancy was due to difficulties explaining the task's purpose, the protocol constraints (no reading), the uneducated participants' inflexible and demanding occupations, and apprehension. In addition, illiterate speakers had limited access to quiet spaces and longer social protocols for foreign visitors.

In Ettimadai, for example, a small vil-



lage with a large engineering school, our local contact and interpreter—a university student who had grown up in a nearby village—recruited (by word of mouth) and arranged short appointments with both literate and illiterate individuals. Participation was on a volunteer basis, as local contacts had suggested. On the first day, the villagers with little or no education all failed to show up for their appointments. The next day, we found that four people had missed the appointment due to work, one had to tend to a sick child, and one reported feeling afraid of what might happen to his voice. When we worked alongside trusted organizations that serve the rural poor (MSSRF and Aravind eye camps), we were much more successful in recruiting and recording villagers, especially illiterate volunteers.

Recording. Illiterate participants in speech studies require novel speech-collection techniques. How do you elicit a particular word from a speaker without saying the word yourself, thus influencing the speaker's word choice and pronunciation? For literate or semiliterate participants, bilingual flashcards with digits written in both numerical and ortho-

graphic form were randomized and shown to speakers one at a time, who were recorded while saying the number aloud. If a participant couldn't read the flashcards, a researcher or interpreter translated the flashcards into a display of fingers (see figure 4). (A fist represented zero.) Unfortunately, we failed to anticipate the cognitive difference between the task of reading numbers aloud and counting fingers,⁴ which has significant effects in the speech domain, such as more disfluencies (filled pauses, false starts, and repetitions) and shouting—for example, “um, thr ... no! FOUR! four.” In combination with environmental conditions, the technique failed to be free of external linguistic influence: recordings often took place outside or in a crowded room, where the task of guessing the number of displayed fingers was often irresistible to “helpful” bystanders.

We had anticipated limited infrastructure and unfamiliarity with technology among rural villagers, however, and built a custom microphone embedded in a telephone handset. The handset was easy for all speakers to use, even those who had seen but never used a phone. It captured quality speech record-

ings in various environmental conditions (average signal-to-noise ratio was 29 dB), avoided the need to fasten equipment to the speaker's clothing, and didn't require the use of a table.

User testing. Another challenge we found in exploring speech technology for illiterate users was in user-testing a spoken dialog system that provides market prices of agricultural commodities and weather. We conducted a Wizard-of-Oz study,⁵ which required one interpreter, one researcher acting as the speech recognizer, and another bilingual researcher observing participants and recording critical incidents. After a short demo by the bilingual researcher, we asked subjects to perform three tasks (for example, "find the weather for today"). The three illiterate participants who agreed to participate performed comparably to literate users.⁶ However, they had more difficulty understanding the nature of a "task" and instead explored the system out of interest or correctly completed unassigned tasks.

In addition, illiteracy appears to affect a speaker's attention to linguistic form. One participant who had never been to school repeatedly used the formal form for "yes" ("amaanga"), which the system didn't recognize. Despite the system's explicit prompts for "amaam," the recognizable form, and multiple instructions from a bystander to say "amaam" instead, she continued to say "amaanga," until finally quitting the system. This experience, along with results from previous user interface studies⁷ and cognitive studies that target illiteracy,⁸ suggest that successful user interface solutions for illiterate users will rely on words and interactions that are meaningful and relevant to everyday language use and won't require illiterate users to memorize icons or command words. Our protocol also included a questionnaire intended to collect linguistic, educational, and regional

information. The participants with very little education completed only one of the three tasks, then politely excused themselves. They didn't complete the questionnaires.

Overall, we learned that individuals with little education can navigate through a dialog system that primarily responds to local terms for "yes" and "no" with very little training. The lengthy (45 minutes) experimental protocol, however, favored literate speakers and left us unable to make quantitative conclusions about differences between the two groups. We also found that traditional data-collection and user techniques favor literate speakers. We're exploring simple, robust speech technologies that adapt to users, thus avoiding the time-consuming, artificial step of recording disjointed instances of speech, which is particularly ill-suited to an illiterate population. Instead, by integrating data collection into a spoken dialog system, we can meet the user's needs (gaining access to relevant information) and the system's needs (gathering speech instances to enable recognition) simultaneously.⁹

Despite these challenges, we plan to continue working in developing regions. The difficulty of the work and the value of the solutions make it quite rewarding. We offer these high-level recommendations on how to approach such work.

- *Plan hard, but be flexible.* Our plans always change in the field because of equipment failures, power or staff problems, or even local crises, and thus agility is probably the most important goal. Nonetheless, detailed planning is worthwhile, exactly because there are so many details: shipping ahead, spare parts and tools, transportation, and time with local partners.
- *Time dilation.* Everything takes longer than expected, which is frustrating given the limited time students have in the field. Often our goals for a trip slip, and we must complete them from abroad or wait until we can return.
- *Bulletproofing.* We spend much more time than usual getting our systems to work reliably and to enable remote management when possible. The distance and lack of local IT staff place a premium on robustness of all kinds, including packaging, theft prevention, and power.
- *Simple UIs.* Similarly, for both end users and IT staff we need simple management UIs, especially for Linux. We try to put all the common management tasks into a simple Web-based UI.
- *Local partners.* We depend on strong, long-term relationships with local partners such as the Aravind Eye Hospital. Our partners provide design input, vendor selection, deployment help, long-term maintenance, and respect within the community. None of this work is possible without them, and they are generally both capable and excited. 

ACKNOWLEDGMENTS

This material is based on work supported by the US National Science Foundation under grant 0326582 and by Intel Corp. Thanks to Alan Mainwaring of Intel Research Berkeley for the power quality measurements and to Michael Rosenblum of UC Berkeley for input on customs and shipping.

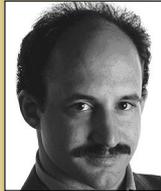
REFERENCES

1. E. Brewer et al., "The Case for Technology in Developing Regions," *Computer*, vol. 38, no. 6, 2005, pp. 25–38.
2. Y. Mugica, *Distributed Solar Energy in Brazil: Fabio Rosa's Approach to Social*

Entrepreneurship; www.nextbillion.net/files/DistributedSolarEnergy_3.pdf.

3. *Provisional Population Totals: Census of India 2001*, Paper 1, Office of the Registrar General, Ministry of Home Affairs, New Delhi, India, 2001; www.censusindia.net.
4. B. Butterworth et al., "Storage and Retrieval of Addition Facts: The Role of Number Comparison," *Quarterly J. Experimental Psychology*, vol. 54A, no. 4, 2001, pp. 1005–1029.
5. J.D. Gould, J. Conti, and T. Hovanyecz, "Composing Letters with a Simulated Listening Typewriter," *Comm. ACM*, vol. 26, no. 4, 1983, pp. 295–308.
6. M. Plauché and M. Prabaker, "Tamil Market: A Spoken Dialog System for Rural India," *Working Papers in Computer-Human Interfaces*, 2006.
7. T. Parikh, G. Kaushik, and A. Chavan, "Design Studies for a Financial Management System for Micro-credit Groups in Rural India," *Proc. ACM Conf. Universal Usability*, ACM Press, 2003.
8. A. Castro-Caldas et al., "The Illiterate Brain," *Brain*, vol. 121, 1998.
9. M. Plauché et al., "Speech Recognition for Illiterate Access to Information," to appear in *Proc. Int'l Conf. Information Technologies and Development*, UC Berkeley Press, 2006.

the AUTHORS



Eric Brewer is a professor in the Computer Science Department of the University of California, Berkeley, as well as director of Intel Research Berkeley. In addition to technology for developing regions, he studies Internet systems and programming languages. He received his PhD in electrical engineering and computer science from MIT. He's a member of the IEEE and the ACM. Contact him at 623 Soda Hall, UC Berkeley, Berkeley, CA 94720-1776; brewer@cs.berkeley.edu.



Michael Demmer is a doctoral candidate at UC Berkeley and an intern at Intel Research Berkeley. His research interests include delay- and disruption-tolerant networking, distributed systems for unusual or challenged network environments, and the application of technology in developing regions. He received his BS from Brown University. Contact him at 74 Chattanooga St., San Francisco, CA 94114; demmer@cs.berkeley.edu.



Melissa Ho is a graduate student in UC Berkeley's School of Information, researching communications and healthcare infrastructure for developing countries. She received her MSc in data communications, networks, and distributed systems from University College London. Contact her at Intel Research Berkeley, 2150 Shattuck Ave., Ste. 1300, Berkeley, CA 94704; mho@sim.s.berkeley.edu.



R.J. Honicky is a doctoral student at UC Berkeley, researching environmental sensors embedded in location-aware cell phones. He received his master's in computer science from UC Santa Cruz. Contact him at 5455 Cory Hall, UC Berkeley, Berkeley, CA 94720-1776; honicky@cs.berkeley.edu.



Joyojeet Pal is a doctoral student in city and regional planning at UC Berkeley, researching community technology initiatives in developing regions. Contact him at 228 Wurster Hall, UC Berkeley, Berkeley, CA 94720-1850; joyojeet@cs.berkeley.edu.



Madelaine Plauché is a postdoctoral researcher of linguistics at the International Computer Science Institute, where she investigates speech tools for limited-resource environments and speech disfluencies. Contact her at ICSI, 1947 Center St., Ste. 600, Berkeley, CA 94704-1198; plauche@icsi.berkeley.edu.



Sonesh Surana is a doctoral student in UC Berkeley's Technology and Infrastructure for Emerging Regions research group. His current work focuses on low-cost networking infrastructure, including using Wi-Fi for long distances and cellphones for rural data collection. He received a BS in computer science from Carnegie Mellon University. Contact him at sonesh@cs.berkeley.edu.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.