

Eric Brewer
 Michael Demmer
 Bowei Du
 Melissa Ho
 Matthew Kam
 Sergiu Nedeuschi
 Joyojeet Pal
 Rabin Patra
 Sonesh Surana
 University of California at Berkeley
 Kevin Fall
 Intel Research Berkeley
<http://tier.cs.berkeley.edu>

The Case for Technology in Developing Regions

I hope the industry will broaden its horizon and bring more of its remarkable dynamism and innovation to the developing world.

—Kofi Annan,
United Nations, 2002

Among the broad set of top-down Millennium Development Goals that the United Nations established in 2000 (<http://www.un.org/millenniumgoals>), one stands out: “Make available the benefits of new technologies—especially information and communications technologies.”

Alongside good governance, technology is considered among the greatest enablers for improved quality of life. However, the majority of its benefits have been concentrated in industrialized nations and therefore limited to a fraction of the world’s population. We believe that technology has a large role to play in developing regions, that “First World” technology to date has been a poor fit in these areas, and that there is thus a need for technology research for developing regions.

Despite the relative infancy of technology studies in developing regions, anecdotal evidence suggests that access to technology has a beneficial economic impact. Cellular telephony is probably the most visible application, but there are many others, some of which we cover in this article.

The World Bank’s infoDev site catalogs hundreds of information and communications technologies (ICT) projects (<http://www.infodev.org>), albeit not all successful. Most of these projects use existing off-the-shelf technology designed for the industrialized world.

Although it is clear that there are large differences in assumptions related to cost, power, and usage, there has been little work on how technology needs in developing regions differ from those of industrialized nations. We argue that Western market forces will continue to meet the needs of developing regions accidentally at best.

ICT RESEARCH FOR UNDERSERVED REGIONS

Evidence from the development of other technologies, such as water pumps and cooking stoves, demonstrates widespread impact from research.^{1,2} Novel ICT has the potential for great impact in a variety of fields ranging from healthcare to education to economic efficiency. However, we do not propose that ICT offers a panacea for the complex problems facing nations on the path to economic development. On the contrary, at best, ICT can enable new solutions only when applied with a broad understanding and a multi-disciplinary approach.




Figure 1.
Schoolchildren in Uttar Pradesh, India, hooking up 220 V AC power by hand to power their PC. They connect and disconnect the power every day—“borrowing” the power from a neighbor.

Any deployment must deal with complex social issues such as underlying gender and ethnic inequalities, as well as existing players on which ICT might have a negative impact. Our strategy on this front is to work closely with social scientists and to partner with strong government or nongovernmental organizations (NGOs), which tend to understand local needs and dynamics in a way that is not possible from afar.

Our research focuses mainly on regions where most of the population makes less than US\$2,000 per year. (Income numbers are normalized for equivalent purchasing power.) The US average income per year is US\$37,800;³ about 65 percent of the world population is below this line, and 80 percent of that number make less than US\$4,500. Progress in India and China in the past few decades has actually improved these numbers overall, but Africa has fallen further behind in per-capita income and thus remains the region that poses the greatest developmental challenges.

Finally, although we have studied ICT in many regions, our direct experience has focused on India, Brazil, Bangladesh, and Uganda, which leads to an inherent bias toward the issues of these countries.

Why now?

Three things make this the right time for deploying ICT in developing regions: the impact of Moore’s law, the increased availability of wireless communications, and the emergence of a more supportive business environment.

The impact of Moore’s law has decreased the cost of computing to fractions of a cent per user. Unfortunately, this cost applies only to the shared infrastructure, not to personal devices. None-

theless, with a focus on sharing, the cost of computing and storage becomes realistic even for the poorest users.

Second, the high-volume production of wireless communications, particularly cellular and Wi-Fi, has decreased their costs as well. For example, the majority of villages in Bangladesh, almost entirely without telephony a decade ago, now have shared cell phones. For rural areas, a wireless infrastructure appears to be the first kind of infrastructure that is affordable. We hypothesize that a successful wireless infrastructure could lead to sufficient increases in rural incomes to make other infrastructure investments viable, such as water and power distribution.

Third, the diffusion of technology worldwide and the growing access to capital have created a favorable environment for entrepreneurship and experimentation, as discussed below. This supportive business environment, combined with the success of franchising as a way to deploy large-scale ICT projects, means that there is a viable path from research to a large-scale impact in the real world.

What should we do?

ICT projects have four main technology needs: connectivity, low-cost devices, appropriate user interfaces (UI), and power.

Rural connectivity is a challenging prerequisite in most ICT projects. Although wireless has broad use in urban areas, most rural areas are without coverage. The low population density in rural areas (even in the US) limits the ability to deploy base stations profitably.

Intermittent—or delay-tolerant—networking requires developing and supporting applications that function without a connected end-to-end networking path. By moving away from the always-on model of traditional voice and IP networks, we can address fundamental disruptions such as intermittent power, extend connectivity via “sneaker” nets, and offer a generally lower-cost infrastructure.

Low-cost devices would enable many applications. The best avenues seem to be enhancing cell phones, enabling sharing, and reducing the basic cost of screens, batteries, and boards.

The generally complex and task-specific nature of modern PC technology creates many UI issues—from basics such as keyboard, font, and multi-language problems to deeper conceptual issues such as the need for speech interfaces for users with limited literacy.

Finally, the availability and quality of electric power form basic hurdles for ICT deployments.

ECONOMIC ISSUES

Economic issues have a significant impact on the opportunities for ICT research in developing regions. Three factors have an impact on providing access to ICT in these areas: investment capital, franchising, and shared technology.

Investment capital

Because it is fundamental to applying innovation, the absence of investment capital hinders advancements in developing regions. In capital-deprived areas, the sole asset for leverage is often family land, which is unfortunately locked up as “dead capital” because there is no way to prove ownership. Inspired by the work of researchers like Hernando De Soto,⁴ many governments are deploying ICT-based land record systems to better utilize these assets.

A revolution in this space has been uncollateralized loans, known as *microcredit*, advanced by Grameen Bank in Bangladesh and FINCA International in Bolivia (<http://www.villagebanking.org>). Grameen Bank, founded in 1975 by Bangladeshi economist Muhammad Yunus,⁵ loaned out an astonishing US\$425 million last year in tiny amounts (US\$50-\$500), generally without collateral, and has a recovery rate approaching 99 percent (<http://www.grameen-info.org/bank/GBGlance.htm>).

In a typical scenario, a Bangladeshi villager uses a microcredit loan for assets such as livestock and then repays the loan with profits from dairy products. These entrepreneurs often expand their business from a single goat to small herds and even become employers—in fact, most loans go to existing businesses. A remarkable 47 percent of loan recipients eventually cross the poverty line, making microcredit one of the most effective tools against poverty.

Franchising

An important outcome of microcredit is that it enables franchising. Traditional franchising—for example, a fast-food business—is a way to scale quickly using the franchisees’ capital. In developing regions, microcredit becomes the source of capital for franchisees. This is uniquely beneficial in areas where the demand for certain services exists, but an effective capital flow does not. Microcredit creates a two-way street: Motivated entrepreneurs with local knowledge have more opportunities, and microcredit lenders can leverage the franchise brand and training for better returns.

The flagship of this model is Grameen Telecom (<http://www.grameen-info.org/grameen/gtelecom>), which grew out of Grameen Bank. Franchisees use

microcredit to buy a cell phone and then operate it as a pay phone for their neighbors. This model has two big benefits: scalability and sustainability.

Using this model, GT now has 95,000 franchisees covering more than 50,000 of the 68,000 villages in Bangladesh, and serving some 60 million people. On average, phone operators earn more than twice the average per capita rural income. GT revenues per rural phone are double that of urban phones.⁶ Since scalability is a fundamental requirement of many projects, we believe the franchising model is a crucial part of many solutions.

Second, franchisees are uniquely good at keeping the system operational. Any ICT project will have unforeseen glitches; the difference here is that someone’s livelihood depends on resolving these glitches. Franchisees form the “human in the loop” that can add a significant level of robustness to any system. This is in stark contrast to most (but not all) donated-computer projects, which have no viable plan for ongoing maintenance and support.

Shared technology

Although personal devices are the norm for developed nations, they are unlikely to succeed in areas where most of the population is below the poverty line. Indeed, most successful projects to date, including Grameen Telecom, depend on shared devices. Given that even cell phones are shared, with a cost approaching US\$40, it is unlikely that personal computers make sense—even a US\$100 laptop would cost too much and use too much power. For example, both the Indian Simputer and the Brazilian Computador Popular did not succeed as personal devices for the poor.⁷ Moreover, Moore’s law will not solve this problem because increasing the number of transistors per chip helps functionality but not cost. CPUs are already a tiny fraction of overall system cost and are dwarfed by the costs of packaging, discrete components, batteries, and screens.

Fortunately, most rural areas have a history of shared technology, including tractors and pay phones. We thus believe that the focus must remain on shared technology, particularly community kiosks, schools, and shared cell phones. There is some room for personal devices, but primarily for specialists such as health workers.

IMPACT AND OPPORTUNITIES

Although environmental and cultural concerns require consideration, the long-term impact of

The focus must remain on shared technology, particularly community kiosks, schools, and shared cell phones.

Adequate healthcare is one of the greatest needs in developing regions.

technology depends on its economic sustainability. The development landscape is littered with failed ICT pilots based purely on short-term charity. Projects can be sustainable either because they serve a public good and have ongoing support—for example, through taxation—or because they are at least self-sustaining financially.

Our research covers both kinds of projects in five areas: health, education, disaster management, e-government, and economic efficiency.

Healthcare

Adequate healthcare is one of the greatest needs in developing regions, which remain home to the vast majority of preventable diseases, including 96 percent of malaria, 95 percent of HIV/AIDS, and 90 percent of tuberculosis (<http://www.theglobalfund.org/en>). Child mortality rates are also high: 10 percent of children under age five die compared with one in 143 in high-income nations (http://www.developmentgoals.org/Child_Mortality.htm).

Although it may be difficult for ICT to have an impact on certain health issues like malnutrition, it can directly impact some areas including disease control, telemedicine, improving doctors' efficiency, offering low-cost diagnostics, improving data collection, and providing patient management tools.

ICT has played a role in some healthcare success stories and other research opportunities are available.

River blindness. ICT has been pivotal in controlling river blindness in West Africa. Hydrological sensors located along 50,000 kilometers of rivers in 11 countries collect data and transmit it via satellite to entomologists, who use forecasting software to compute the optimal time, dosage, and area over which to spray larvicide aerially to destroy the disease-spreading blackfly.⁸ This endeavor protects 30 million people and has also reclaimed 100,000 square miles of farmland.⁹

Child mortality. At a cost of less than US\$2 per capita, the Tanzania Essential Health Initiatives Program conducted periodic computer analyses of prevailing cause-of-death data collected from rural households by health workers in two Tanzanian districts.¹⁰ Results indicated that only a fraction of the budget was being spent on the actual killer diseases, which led to acute shortages of key medicines. With 80 percent of the people dying at home and leaving no records, this insight was previously hidden. Strategically retooling the budget in response to

timely reports on prevailing disease characteristics led to more effective care and dramatically reduced child mortality by an average of 40 percent in five years—from 16 percent to less than 9 percent for children under age five.

Aravind Eye Hospitals. The mission of this network of five self-sustaining hospitals in Tamil Nadu, India, is to sustainably eradicate blindness. Aravind, which treated nearly two million patients and performed more than 200,000 eye surgeries in 2004,¹¹ emphasizes using existing technology, even locally manufacturing its own intraocular lenses. A cataract surgery costs only about US\$10; although about half the patients are treated for free, the practice is still profitable. The network uses a VSAT satellite connection at 384 Kbps to videoconference among hospitals and also to communicate with a mobile van conducting “eye camps” in rural areas.

Our experience working with Aravind has revealed several research opportunities.

Telemedicine. Existing telephony networks do not provide adequate support for videoconferencing with centrally located doctors to enable remote consultation and health worker training. Telemedicine—using telecommunications for the remote diagnosis and treatment of patients—requires designing low-cost, low-power, long-range, and high-bandwidth wireless technology.

Computer-assisted diagnosis. There is a growing shortage of health workers, with the poorest regions hurt most (<http://www.globalhealthtrust.org>). To address this problem, in addition to other recommendations, there is an urgent need to automate simple diagnostic tasks. Software such as image-recognition tools or decision-support systems and hardware such as the ImmunoSensor chip (<http://www.coe.berkeley.edu/labnotes/1003/boser.html>) can reduce the strain on overburdened health staff, improve efficiency, and mitigate the lack of infrastructure such as labs.

We are currently testing a prototype that diagnoses diabetic retinopathy from digital retina images. Without automation, doctors spend about five minutes per diagnosis, even though the surgery itself only takes 15 minutes. Assuming 90 percent accuracy, which we are approaching, partially automated diagnosis would result in an annual savings of thousands of hours per hospital, which could dramatically increase the number of patients receiving surgery.

Field-worker support. With low utilization of healthcare services and poor quality of health data, many organizations are sending semitrained health workers into rural areas. These workers typically

are equipped with handheld devices to reduce errors and time for data transcription.

Although pilot projects show promising results (<http://www.healthnet.org/usersurvey.php>; <http://www.healthnet.org/coststudy.php>), there are problems due to frequent power outages (<http://www.opt-init.org/framework.html>) and poor telecommunications. This is a strong motivator for work in developing low-cost intermittent networking devices.

Education

According to a UNESCO study of global education, only 37 of 155 countries have achieved a universal primary education.¹² In most high-income countries, students are expected to attend school for at least 11 years, but only two out of the 37 low-income countries offer the same level of education.

In recent years, several attempts have been made to integrate ICT into rural and low-income urban schooling. By combining technology with sound educational principles and teaching practices, many of these initiatives have demonstrated increased learning.

Azim Premji Foundation. Since 2001, electronic courseware has been developed in local languages to aid classroom learning in rural schools in the states of Karnataka and Andhra Pradesh, India. Students use shared computers and work in small teams to run the courseware, and the content is designed so that students work through the material in a self-paced manner and via discussions with one another.

A randomized experiment involving 2,933 students from 32 schools in both states showed statistically significant improvements on learning achievement tests for students in Andhra Pradesh but not in Karnataka (<http://www.azimpremjifoundation.org/downloads/ImpactofCALonLearningAchievements.pdf>). One explanation offered by the investigators is that teachers were not part of the computer-aided learning process in the Karnataka pilot. In Andhra Pradesh, teachers accompanied the students to the telecenters, but students in Karnataka were presumably responsible for using the courseware by themselves. As of March 2005, this initiative covered more than 10,200 schools.

Digdarshan. In 2000, researchers placed a solar-powered computer in a rural school in Uttar Pradesh, India, and developed Hindi CD-ROM courseware for it. Students used the computer in small groups to help one another understand the material, and enjoyed the interactive quizzes to the

extent that they persisted in retaking the exams until they attained full scores; in the process, they helped one another to understand why their initial answers were incorrect. More importantly, given the shortage of qualified teachers, the courseware fills in gaps in the teachers' knowledge and aids senior students in coaching their juniors. Finally, parents became more willing to send their children to school because they were aware of the importance of IT literacy.

As of April 2005, Digdarshan has reached 700 schools with 700 more in progress.

Based on these examples, we have identified several research opportunities in this area.

Digital story authoring tools. We found that children are highly motivated to create multimedia digital stories that impress fellow students. Creating these stories fosters active learning because the authors must explain academic concepts while they are developing their writing and communication skills. However, current authoring tools such as Microsoft PowerPoint are too complicated, and simpler tools like KidPix are not necessarily culturally appropriate.

Tools can be localized by evoking cultural icons and mythology such as the Amar Chitra Katha comics in India. For example, with the Suchik project research prototype (<http://www.iitk.ac.in/MLAsia/suchik.htm>), students create comic stories with user-defined characters. We envision authoring tools that children can use to create interactive stories and games.

Local content repository. To the extent that local content exists, its limited availability often restricts its impact. We envision a Web-accessible repository where students and teachers can store and retrieve the digital stories, games, and other electronic content that they author. Access to an audience via such a repository is expected to motivate students to create high-quality work. This repository can also act as a digital library of resources, such as background articles, raw data, and clip art that students can use when researching and authoring content. Finally, this repository can manage content created by experienced teachers or by the well-educated diaspora living abroad, which other teachers can use and even customize. The project requires only intermittent networking for each school.

Disaster management

The Indian Ocean tsunami disaster in December 2004 was a tragic reminder of the urgent need for better disaster warning and relief systems. ICT has

Frequent power outages and poor telecommunications are strong motivators for work in developing low-cost intermittent networking devices.

Software is a key component for addressing the problem of matching donations to local needs.

a major role to play in both areas, but had little impact when the tsunami occurred. The focus in the media has been on tsunami detection, but an effective disaster notification mechanism, such as the early warning system for cyclones in Bangladesh,¹³ is arguably more important and is an application in which ICT can directly save lives.

Small victories. The MS Swaminathan Research Foundation runs a communications network in rural areas of Pondicherry in southern India through a web of “information villages” networked through wireless connections. The MSSRF network normally is used to provide communications, weather forecasts, wave and fish location patterns, and other similar services to coastal and inland villages.

On the day of the tsunami, Nallavadu, a village of 3,600 residents, learned of the tsunami via phone and broadcast a warning over a public address system, and no lives were lost. A second village, Veerampattinam, used a similar broadcast after the first wave hit, and lost only one life out of 6,300. Later, these two villages used the databases in their knowledge centers to organize relief measures and distribute aid.

MSSRF is now endeavoring to install similar systems and knowledge centers in neighboring coastal villages (http://www.mssrf.org/notice_board/announcements/tsunami/tidal_tragedy.htm).

Fast-deployment networks. Rural wireless coverage becomes particularly important after a disaster. Although cellular coverage would be ideal, even intermittent networking using a range of technologies, including “sneaker net” and satellite communications, would be helpful. For example, the Grameen Telecom village cell phones were extremely helpful in 2004 when the worst flood in modern history covered 60 percent of Bangladesh and left 30 million homeless.¹⁴ However, many disaster-prone areas are too rural to have existing cellular coverage.

Relief logistics. Matching donations to local needs is consistently a problem. Following the recent tsunami, there were several instances of wasted donations. For example, donors kept sending clothes, while disaster areas did not receive much-needed rice and shoes.

Software is already a key component for addressing these problems. One example is the Fritz Institute’s Humanitarian Logistics Software, a Web-based application for automating the logistics behind mobilization, procurement, distribution, and finance processes (<http://www.fritzinstitute.org>).

Another organization, Auroville, is addressing these issues independently by establishing Tsunami Rehabilitation Knowledge Centers with the express purpose of facilitating the flow of information between organizations and villages and providing better coordination among the NGOs (<http://www.auroville.org/tsunami/projects.htm>). Relatively simple shared Web sites, even with intermittent connectivity, are central to the strategy for this collaborative effort.

E-government

The case for technology in e-governance encompasses three broad application areas: public information, digitization of records, and transactions involving the state. There is evidence that e-government services offer significant benefits to users¹⁵ and that developing countries are embracing the Internet as a citizen-government interface.¹⁶ However, many areas—including sub-Saharan nations and parts of Asia—lack both online services and basic computerization of processes within the government.

Friends network and Akshaya. A study of the Fast Reliable Instant Efficient Network for Disbursement of Services bill-payment network (<http://www.keralaitmission.org/content/egovernance>), which offers one window that replaces more than 1,000 forms and service requests at district headquarters, shows that even in the most-rural areas more than 95 percent of people prefer online transactions. The Friends study indicated that single-counter e-government kiosks could handle an average of 400 transactions daily, whereas a human operator could perform only a fraction of that service.

The Friends network represents savings not only to the government agency deploying the service but also directly to users, both in terms of money and time saved per month. There is also evidence that e-government services, even when run by private entrepreneurs on behalf of the government, can still be sustainable (<http://www.drishtee.com>).

The Akshaya project (<http://www.akshaya.net>) in northern Kerala extended the Friends program by creating a Web-based interface. Currently, more than 100 e-governance kiosks are providing e-payment services. These online services are established with the assistance of local government bodies and are run by entrepreneurs who charge small fees for paying bills. The entrepreneurs in turn maintain accounts with the state and pay online after accepting cash payments.

Concurrent with this program, the state trained at least one person in each household in basic com-

puter usage with the aim of eventually enabling people to use e-government services on their own. The e-pay services are among the most successful in the state, and several entrepreneurs have raised their credit limits to handle the high demand.¹⁷

Bhoomi. The Bhoomi land records project (<http://www.bhoomi.kar.nic.in>) in Karnataka, India deals with the problem of consolidation of multiple forms of land records, accumulated over decades or even centuries, and representing different systems of land records under various regimes throughout history. Bhoomi seeks to digitize these records into one consistent format.

The state government digitized more than 20 million land records held by 6.7 million citizens in more than 2,700 villages in the state. The digitization included visits by land surveyors, several levels of checks and balances to ensure veracity of land claims, and the creation of a fingerprint authentication system that integrates all current transactions by village accountants. A secure database at the back end is meant to ensure that the land data thus gathered remains reasonably protected from the machinations of local officials, who often are accused of harassment and extortion of marginal villagers.

A touch screen for land record verification is kept in a public place to ensure that villagers can check on the status of their land records. Printing records takes about two minutes and avoids a complex application system that previously took up to 30 days.

Bhoomi has some important long-term prospects. The state government currently is converting the data into a geographic information system format. Banks already use Bhoomi data for farm credit, a process that currently takes fewer than five days compared to one month in the past. Banks also use the records to estimate farm credit requirements and asset leverage. The previous shortage of such information was an important factor that limited access to capital in rural areas. Finally, land-use planners can use the system to get a better idea of crop patterns.

Economic efficiency

Economic efficiency depends on information and communication. Without an ICT infrastructure, communication relies on physical contacts that are expensive in terms of both time and money. As in industrialized nations, communications networks can lead to more efficient and transparent markets. ICT can facilitate improvements in communications, fair-market price discovery, and supply-chain management.

Grameen Telecom. Spawned by Grameen Bank, GT's goal is to provide affordable telecommuni-

cations access to rural villages in Bangladesh. The basic GT village phone model consists of franchisees who offer telephony using a cell phone and charger leased from the company. Each franchisee earns revenue by reselling phone minutes, which they buy from GT at a flat rate. Wealthier urban customers subsidize the base stations.

Surveys show that 50 percent of GT phone use is economic in nature.¹⁴ Survey responses also indicated that in half the cases, the only alternative form of communication was via travel or courier, which was two to eight times more expensive and required 2.7 hours on average.¹⁸

Enhanced communications access has improved economic transparency in the GT phone villages. For example, a chicken egg farmer found that by calling the market for the fair price of eggs, she can negotiate with middlemen for a better offer. Before she had access to the phone, the farmer had little negotiating power.¹⁹

In addition to market pricing, villages also use the phone to guarantee full delivery of remittances from relatives working abroad, which accounted for 4.5 percent of Bangladesh's GDP in 2001.²⁰ The villagers use the phone in two ways. First, relatives abroad can call and verify that the correct amount of money was received. Second, by using the phone to obtain foreign exchange rates, the villagers can negotiate fair rates.

GT benefited to some extent from the high population density of Bangladesh, which enables profitable base stations. Lower-density areas, which include most rural regions, require new wireless technology. In addition, many of these economic activities require only intermittent networking.

e-Choupal. ITC's e-Choupal initiative (<http://www.echoupal.com>) is a computer kiosk deployed to create an efficient agricultural supply chain. Traditionally, farmers sold their crops to ITC via bidders at *mandis*—aggregated market yards. This scheme had several inherent inefficiencies: insufficient knowledge of current market prices, no differentiation of price with respect to crop quality, and high transportation and handling costs.

In 2000, ITC began building information kiosks in the Indian state of Madhya Pradesh to provide farmers a direct market for their crops. Each kiosk is equipped with a PC, a printer, and power backup, and the network connection is provided via phone or VSAT.²¹ The kiosks also have automatic moisture analyzers that measure the quality of crop samples. ITC subsidizes the equipment, and a local farmer houses the kiosk.

The only alternative form of communication was via travel or courier, which was two to eight times more expensive.

A consistent theme in ICT research is the need for network infrastructure.

The farmers use the kiosks for informational and commercial tasks. The e-Choupal Web portal provides farmers with market prices, weather predictions, and educational material on agricultural practices. Because the quality of the crop can be measured locally, the farmers immediately receive payment for their crops from the kiosk. Retail companies also can use the kiosks to sell their agricultural products directly to the farmers. As a side effect, NGOs have used the networking infrastructure to pursue development work.²¹

The e-Choupal project has expanded from Madhya Pradesh to provide a total of 4,200 kiosks covering 29,500 villages in six Indian states. An open question is how to leverage kiosks driven by supply chains to provide a broader array of services.

EARLY RESEARCH AGENDA

The presence of a communication and power infrastructure, appropriate user interfaces, and access to inexpensive computing devices appear to be areas ripe for innovation.

Networking infrastructure

A consistent theme in ICT research is the need for network infrastructure. Unfortunately, most current networking technologies are still not available for developing countries. For example, the typical cost for a telephone landline is US\$400, which is acceptable in the US, where 90 percent of households can afford to pay \$30 a month for telephone service. In contrast, in India, more than 60 percent of the population can afford at most US\$5 a month for communications.

Moreover, communication technologies are even less practical for rural areas. In a typical network, more than 70 percent of the cost is in the access network, not in the backbone. This constitutes an important limitation for rural deployment, since user density and consumer paying capacity are low. Thus, the recent growth of cellular operators has been an urban phenomenon. Even though the majority of the developing world population lives in rural areas—for example, 74 percent in India (<http://www.censusindia.net>)—most rural areas remain uncovered (<http://www.auroville.org>).

Rural wireless. To help alleviate this situation, we have begun to research cost-effective technologies for access networks in sparsely populated rural areas. Among the available technologies, wireline networks are unaffordable. Although satellite-based systems such as VSAT can cover remote areas, they are prohibitive in terms of initial and recurring costs;

cellular technology is designed for high population densities and typically does not have long range.

Three technologies look promising: modified Wi-Fi, CDMA450, and 802.16.

We believe that the most promising solution to extend coverage to rural areas is a mixture of point-to-point and point-to-multipoint wireless technologies. These technologies could form a backbone that provides connectivity down to the village level, combined with a Wi-Fi hot spot or a cellular base station for access to individual households.

Our technological focus has primarily been on adapting the IEEE 802.11 (Wi-Fi) family of technologies (<http://grouper.ieee.org/groups/802/11>) to provide the backbone network. The 802.11 standard enjoys widespread acceptance, has huge production volumes, with chipsets costing as low as US\$5, and can deliver bandwidth of up to 54 Mbps. Furthermore, point-to-point links using 802.11 high-gain directional antennas can span impressive distances up to 88.7 km (<http://www.wired.com/news/culture/0,1284,64440,00.html>).

Despite its attractive features, the 802.11 Media Access Control protocol was designed primarily for short-distance, broadcast media with tens of mobile hosts contending for access. Consequently, it is inefficient for long-distance connections among a small set of known hosts. However, small modifications at the MAC layer should overcome these problems, and this is an area of active research.^{22,23} With these modifications, we believe that Wi-Fi is the most promising short-term alternative.

In contrast, a number of existing deployments have used proprietary technologies to provide village-level connectivity. The Akshaya network uses a combination of Wi-LAN (<http://www.wi-lan.com>) and wireless in local loop (WipLL, <http://www.airspan.com>) links. Unfortunately, the production volumes for these technologies are too low to enable costs comparable to mass-produced and standardized Wi-Fi. Tenet's CorDECT solution (<http://www.midascomm.com>) has a low cost due to mass production, but it only delivers a 70-Kbps peak data rate.

The CDMA450 (<http://www.cdg.org/technology/3g/cdma450.asp>) cellular technology operates in a lower frequency band (450 MHz) than traditional Global System for Mobile (GSM) communication and code division multiple access (CDMA) solutions, achieving larger coverage per base station and thus lower cost (a radius of 49 km versus 35 km using standard GSM). Early successes include China and Romania.

Another promising alternative is Wi-Max, the IEEE 802.16 standard for wireless metropolitan-

area networks (<http://grouper.ieee.org/groups/802/16>). From a technical point of view, the standard is well suited to the problem at hand: It accommodates thousands of users, uses the available spectrum efficiently, covers ranges of more than 50 km, and has quality-of-service features. Unfortunately, at this time the standard is not finalized, and prestandard equipment sells at prohibitive prices. However, if WiMax shares the market success of Wi-Fi, it will likely evolve into a useful tool for rural networking infrastructure.

One advantage of CDMA450 and 802.16 is that they require fewer large towers than a Wi-Fi mesh, thus reducing the overall cost for flat terrain.

Early results. To demonstrate the efficacy of Wi-Fi for rural connectivity, we have undertaken several deployments in India in addition to our wireless testbed in Berkeley. In one such link deployed for the Aravind Eye Hospital, we have demonstrated a bandwidth of 4 Mbps over a 10.5-km distance. This is a vast improvement over the 33.3 Kbps achieved by existing WiLL links.

To address the inefficiency of Wi-Fi in a multihop network with long-range, point-to-point links, our group has identified MAC layer modifications that significantly increase bandwidth efficiency and minimize interference when multiple radios are installed on the same tower. Our approach changes the contention-based CDMA scheme to a time division multiple access scheme, yet still operates with standard, inexpensive 802.11 hardware. We are currently testing this approach on our wireless testbed in Berkeley and are planning to use it for future deployments of larger multihop networks in India.

Intermittent networking. The focus of nearly all networking technology—including telephony, the Internet, and most radios—is on providing continuous, real-time, synchronous communication. However, the connectivity and power infrastructure required to support such continuous access may be too costly to justify deployment in many developing regions.

We believe that many useful applications could instead be designed around asynchronous communication and only intermittent connectivity and that this form of infrastructure can be created at a significantly reduced cost. Some concrete examples of such applications include electronic form filling (much like postal forms), crop or commodity price updates, weather forecasts, offline search engine queries, and more traditional applications such as e-mail and voicemail.

To this end, we have been continuing the development of a *delay-tolerant networking* architec-



ture.²⁴ DTN provides reliable delivery of asynchronous messages (like e-mail) across highly heterogeneous networks that can suffer from frequent disruption. Applications that are developed for use with DTN do not make assumptions about the timeliness of network transactions, so DTN messages can be delayed or batched to reduce cost or power.

DTN can leverage several well-known networking technologies, but it also can take advantage of some nontraditional—and perhaps nonobvious—means of communication. One example is *data mules*, which are devices used to ferry messages from one location to another in a store-and-forward manner.

The DakNet project deploys data mules in the form of bicycles or motorcycles that make regular trips to remote villages.²⁵ In many developing regions, rural bus routes or mail carriers regularly visit villages and towns that have no existing network connectivity. By applying DTN technologies, we can leverage the local transportation infrastructure as a novel networking infrastructure, providing connectivity (albeit intermittent) at a significantly reduced cost.

Accommodating networks such as these that can be partitioned regularly and that can utilize highly varying underlying technologies (such as data mules) requires new approaches to network routing and naming and protocol design. Many routing

Figure 2. Students and village residents putting up a long-distance Wi-Fi link at sunset in southern India. Wireless technologies have the potential to provide a backbone that offers connectivity down to the village level.

A delay-tolerant network system can deliver data reliably even in the face of a wide range of disruptions.

algorithms typically operate on a routing graph with one large, connected component and fail to function well when the graph is frequently partitioned, even when the outages are completely predictable. However, by taking advantage of in-network storage and mobile routers and using novel approaches to message routing, a DTN system can deliver data reliably even in the face of a wide range of disruptions.

Early results. In collaboration with others in the Delay Tolerant Networking Research Group (<http://www.dtnrg.org>), we have extended the DTN architecture²⁴ and have proposed novel protocols for message transmission and forwarding. We have built a functional reference implementation and demonstrated its operation on a range of platforms from PCs to PDAs.²⁶ We have also taken some theoretical steps toward developing routing algorithms that take into account both predicted and unexpected outages.^{27,28} Finally, we are looking at storage systems that provide availability and best-effort consistency on top of intermittent networks.

Inexpensive computing devices

Most previous attempts to develop inexpensive computing devices were commercial failures.⁷ Looking at these projects, we can identify a number of factors that need to be understood when designing low-cost devices. In addition to determining whether to follow a model of shared or personal device ownership, researchers must decide whether to design general-purpose devices (like PCs) or devices tailored for specific tasks such as voice mail, form filling, or data collection. Task-specific devices, such as cell phones, require less training and can be less expensive (via omission) than more general-purpose devices, yet they still can have sufficient volume for economies of scale.

A third question concerns the correct form factor for the device: At one end of the spectrum are PCs and smaller variants like AMD's 50x15 personal Internet communicator device (<http://50x15.amd.com/home/default.aspx>) that consume more power, while at the other end are handheld devices like cell phones. We are pursuing the design of smart phones as a way to increase the impact of cell phones.

Other important research topics that are relevant to reducing the cost of computing devices are low-cost displays and batteries and specialized sensors for testing water, air, or soil quality and for disease detection.

User interfaces

Although UI design and the field of human-computer interaction has made much progress, even the basic components of computing interfaces encounter problems in developing regions. Mouse motions and clicks are not intuitive to the inexperienced user, and differences in language and alphabet render a single keyboard much less useful.

Despite some success with Unicode, the representation of most languages remains in progress, and there are still limited resources for localization of existing content and software. Even given representation and localization, text-based interactions with computers render these devices useless for illiterate or semiliterate users. Finally, the cost of the display is a significant component of the total cost of most computing devices.

All of these issues imply significant ICT research challenges. We also expect DTN-based systems to present usability challenges due to their intermittent online status. This hypothesis is based on fieldwork in Uganda,²⁹ in which the microfinance technology that we evaluated had support for both online and offline transactions. Several users who were newly introduced to this system seemed to encounter difficulty understanding disconnected operations and how they differed from real-time transactions, which appeared to be more intuitive.

Speech. One avenue is finding mechanisms for real-time speech recognition by low-cost, power-constrained devices. We believe this problem can be made tractable by dynamically switching among application contexts, thus limiting the number of words that are considered for recognition at any point in time, without limiting the total number of usable words.

In our approach, centralized computers perform the computationally intensive tasks of speech modeling and training offline, leaving only the task of actual recognition for the simplified devices. Furthermore, we believe that many useful applications in the developing world can be constructed around a limited-vocabulary voice-based interface, similar to VoiceXML. Concrete examples of voice-amenable applications include disease diagnostic devices and commodity price distribution.

Early results. As part of our fieldwork, we gathered samples of spoken numbers and simple commands from Tamil speakers in India and Berkeley. Initial results showed that we can achieve successful speech recognition training using relatively few speakers (about 30). We also completed the design for a low-power hardware-based speech recognizer chip.³⁰ We estimate a 100 percent duty cycle power

dissipation of around 19 mW, which is orders of magnitude lower than existing solutions.

Our simulations demonstrate accuracies of up to 97 percent for speaker-independent recognition in both English and Tamil, which we believe to be sufficient for UI recognition. The estimated size of 2.5 mm² per chip in 0.18- μ m technology confirms the potentially very low production cost.

We are also pursuing the UI design for educational software, combining the issues of literacy, usability, learning, and motivation. Our initial fieldwork in India suggests that even though the UI must be sufficiently simple and attractive for children to use, it must also balance fun and learning. More importantly, the UI should promote the child's personal development. For example, the results of our in-class experiment with a note-taking application suggest that UIs could be designed as scaffolds that develop more effective note-taking habits among students, as well as direct the user's attention to critical topics during a lecture.³¹ Learner-centered design seems especially critical when there is a lack of qualified teachers.³²

Power systems

Although there continues to be much effort in rural electrification, the lack of power remains a fundamental obstacle to the adoption of ICT. Despite some success with providing solar power in India and Kenya, grid coverage is rare in rural areas. Even where electricity is available, we have found the low quality of power to be a huge practical problem for ICT. Note that lighting and cooking—the primary uses—are relatively immune to this limitation.

Initial measurements with high-quality data loggers show huge variation in voltage and current, including short spikes of up to 1,000 V (for a 220 V AC system), long brownouts at 150 V, and frequent outages, both long and short. Using voltage stabilizers or UPS systems is an expensive solution that adds significantly to the real cost of ICT. We expect that changes in PC motherboards or power supplies would be a more effective strategy.

We have also found great inefficiency in the solar systems used for ICT. A typical system uses solar panels and 12-V lead-acid batteries, an inverter to bring the voltage up to 220 V AC, and then PC power supplies to bring the voltage back down to 12 V DC. The standard power supply must generate ± 12 V, ± 5 V, and ± 3.3 V, all at worst-case currents; this requires a fundamentally expensive power supply.

Batteries are also a problem. Existing systems abuse the batteries, which leads to a very short life. The lead in lead-acid batteries has a huge environmental impact, and the common practice of adding acid to old batteries leads to health concerns. In the long term, we must replace lead-acid batteries with a more sustainable option: A built-in smart-charge controller seems a likely contributor to the solution.

The lack of power remains a fundamental obstacle to the adoption of ICT.

PRAGMATIC CHALLENGES

Numerous pragmatic issues have an effect on ICT projects, including design and deployment strategies, transition planning, and the use of open source software.

Codesign, codeploy

Field deployments repeatedly suggest that ICT projects for underserved populations pose unique design challenges. The design of the Simputer (<http://www.simputer.org>), one of the most prominent device projects in this space, had both a needs assessment and a UI design component, yet studies have suggested that the outcome was still hindered by a mismatch between technologists and users.⁷ In fact, the Simputer eventually became a PDA without development goals.

Given that unique regional and cultural characteristics can play a large role in determining a project's success, effective codesign requires using local knowledge to understand the appropriateness of certain technologies over others. Similarly, there is a strong case for coupling the use of local knowledge for design with partnerships in deployment.

ICT researchers also benefit from the stable long-term agendas of NGOs compared with governments. Often the only organized bodies that reach remote populations are NGOs and local government bodies such as village councils. NGOs also tend to be easier to access and relatively transparent due to outside funding. However, we have found that motivating these groups requires building relationships and showing concrete early results.

Transition planning

Experience suggests that a subtle issue for ICT projects is the need for a careful transition plan. For example, in ITC's e-Choupal project, the kiosk system essentially disintermediates agricultural middlemen. Such situations are not at all uncommon—e-governance operations often depend upon corrupt officials supporting systems that disen-

A subtle issue for ICT projects is the need for a careful transition plan.

franchise their “gray” livelihoods. To get the support of such intermediaries, the transition plan must include their interests. Thus, agricultural kiosk projects might consider either hiring former middlemen as kiosk operators or simply paying them off. The failure to make a gentle transition can indeed cripple ICT projects that otherwise seem likely to succeed.³³

Similarly, issues such as land record titling, widely considered as a key path to progress, present transitional challenges that could derail or even exacerbate developmental gaps. Examples from southern India and Cambodia present some negative impacts of technological innovation, when modern land-titling technologies caused destitution among the unprotected poor, who often have lived for years as squatters on government land (<http://www.slate.com/Default.aspx?id=2112792>).

Being mindful of the informal economies that innovative technologies disrupt is essential, and the appropriate transition mechanisms must be incorporated to protect the interests of those with the least voice in the creation or deployment of these systems.

Open source software

In general, developing regions have a default preference for open source software on the premise that it is free. In practice, however, “paid” software such as Microsoft Windows is also free due to piracy, and the legitimacy of software becomes an issue only when the funding source requires it. The more persuasive argument for open source is the ability to localize and customize. Governments in several African, Asian, and Latin American nations have funded efforts for customized local versions of software. This is a critical issue for developing regions where technological solutions often must be microlocalized for markets that are too small to attract the interest of large producers. We have also found serious problems with spyware and viruses in ICT pilots that using open source software could avoid—at least in the short run.

That being said, there are important advantages to Windows. Due to its ubiquitous use, Windows knowledge is viewed as a valuable skill. In Brazil, Windows vocational training for the poor appears to be a viable business.³⁴ More people are trained in using Windows, which can facilitate system administration. Our work at Berkeley is all open source, but our partners use both kinds of software and we expect this will continue.

Our objective is to convince researchers that ICT can play a large role in addressing the challenges of developing regions and that there is a real need for innovative research. The needs of developing regions are both great and unique, and thus lead to a rich and diverse research agenda. Moreover, as these needs are different from those of industrialized nations, market forces will continue to meet them at best accidentally. Because the needs are great, we must do better. Providing ICT for developing regions is not easy, but it is uniquely rewarding. We encourage the ICT research community to take on the challenge. ■

References

1. P. Polak, B. Nanes, and D. Adhikari, “A Low-Cost Drip Irrigation System for Small Farmers in Developing Countries,” *J. Am. Water Resources Assoc.*, Feb. 1997.
2. K.R. Smith et al., “One Hundred Million Improved Cook Stoves in China: How Was It Done?” *World Development*, vol. 21, no. 6, 1993, pp. 941-961.
3. B. Milanovic, “True World Income Distribution, 1988 and 1993: First Calculation Based on Household Surveys Alone,” *The Economic J.*, Jan. 2002, pp. 51-92.
4. H. De Soto, *The Mystery of Capital: Why Capitalism Triumphs in the West and Fails Everywhere Else*, Basic Books, 2000.
5. M. Yunus, “The Grameen Bank,” *Scientific American*, Nov. 1999, pp. 114-119.
6. N. Cohen, “What Works: Grameen Telecom’s Village Phones,” World Resources Inst., 2001, <http://www.digitaldividend.org/pdf/grameen.pdf>.
7. R. Fonseca and J. Pal, “Bringing Devices to the Masses: A Comparative Study of the Brazilian Computador Popular and the Indian Simputer,” *Proc. South Asia Conf.*, UC Berkeley: Trends in Computing for Human Development in India, Feb. 2005.
8. E. Servat et al., “Satellite Data Transmission and Hydrological Forecasting in the Fight against Onchocerciasis in West Africa,” *J. Hydrology*, vol. 117, 1990, pp.187-198.
9. World Health Organization, *Success in Africa: The Onchocerciasis Control Programme in West Africa, 1974-2002*, 2003.
10. D.D. Savigny et al., *Fixing Health Systems: Linking Research, Development, Systems, and Partnerships*, Int’l Development Research Centre, 2004.
11. Aravind Eye Hospitals, *Ann. Activities Report*, 2004.
12. UNESCO Institute for Statistics, *Global Education Digest 2004: Comparing Education Statistics Across the World*, UNESCO, 2004.

13. M.H. Akhand, "Disaster Management and Cyclone Warning System in Bangladesh," *Early Warning Systems for Natural Disaster Reduction*, J. Zschau and A.N. Koppers, eds., Springer, 2003.
14. A. Bayes, J. von Braun, and R. Akhter, "Village Pay Phones and Poverty Reduction: Insights from a Grameen Bank Initiative in Bangladesh," *Discussion Papers on Development Policy*, ZEF Bonn Center for Development Research, June 1999.
15. Dutch Ministry of the Interior and Kingdom Relations, "Capgemini Netherlands and TNO. Does e-Government Pay Off?" Nov. 2004.
16. P. Banerjee and P. Chau, "An Evaluative Framework for Analyzing e-Government Convergence Capability in Developing Countries," *Electronic Government*, vol. 1, no. 1, 2004, pp. 29-48.
17. S. Nedeveschi et al., "A Multidisciplinary Approach to Studying Village Internet Kiosk Initiatives: The Case of Akshaya," *Proc. Policy Options and Models for Bridging Digital Divides*, Global Challenges of e-Development; <http://www.globaledvelopment.org/forthcoming.htm>.
18. D.D. Richardson, R. Ramirez, and M. Haq, *Grameen Telecom's Village Phone Programme in Rural Bangladesh: A Multi-Media Case Study*, Telecoms Development Group, Mar. 2000.
19. S. Bhatnagar and A. Dewan, "Grameen Telecom: The Village Phone Program; http://poverty2.forumone.com/files/14648_Grameen-web.pdf.
20. International Monetary Fund, *Balance of Payments: Statistics Yearbook*, 2003.
21. D.M. Upton and V.A. Fuller, "The ICT e-Choupal Initiative," Harvard Business School Case N9-604-016, Jan. 2004.
22. P. Bhagwat, B. Raman, and D. Singh, "Turning 802.11 Inside Out," *Proc. Hot Topics in Networks (HotNets-II)*, ACM Digital Library, 2003.
23. B. Raman and K. Chebrolu, "Revisiting MAC Design for an 802.11-Based Mesh Network," *Proc. Hot Topics in Networks (HotNets-III)*, ACM Digital Library, 2004.
24. K. Fall, "A Delay-Tolerant Network Architecture for Challenged Internets," *Proc. ACM SIGCOMM Conf.*, ACM Press, 2003, pp. 145-158.
25. A.S. Pentland, R. Fletcher, and A. Hasson, "DakNet: Rethinking Connectivity in Developing Nations," *Computer*, Jan. 2004, pp. 78-83.
26. M. Demmer et al., *Implementing Delay-Tolerant Networking*, tech. report IRBTR-04-020, Intel Research Berkeley, 2004.
27. S. Jain et al., "Using Redundancy to Cope with Failures in a Delay-Tolerant Network," to appear in *Proc of ACM SIGCOMM Conf.*, ACM Press, Aug. 2005.
28. S. Jain, K. Fall, and R. Patra, "Routing in a Delay Tolerant Network," *Proc. ACM SIGCOMM Conf.*, ACM Press, Aug. 2004, pp. 27-34.
29. M. Kam and T. Tran, "Lessons from Deploying the Remote Transaction System with Three Microfinance Institutions in Uganda," <http://www.csberkeley.edu/~mattkam/publications/UNIDO2005.pdf>.
30. S. Nedeveschi, R. Patra, and E. Brewer, "Hardware Speech Recognition on Low-Cost and Low-Power Devices," *Proc. Design and Automation Conf.*, 2005.
31. M. Kam et al., "Livenotes: A System for Cooperative and Augmented Note-Taking in Lectures," *Proc. ACM Conf. Human Factors in Computing Systems (CHI 05)*, ACM Press, 2005, pp. 531-540.
32. C. Quintana et al., "Learner-Centered Design: Reflections and New Directions," *Human-Computer Interaction in the New Millennium*, J. Carroll, ed., Addison-Wesley Professional, 2001.
33. R. Kumar, "e-Governance: Drishtee's Soochana Kendras in Haryana, India," *Proc. South Asia Conference: Trends in Computing for Human Development in India*, 2005.
34. C. Ferraz et al., "Computing for Social Inclusion in Brazil: A Study of the CDI and Other Initiatives," *Bridging the Divide 2004*; http://bridge.berkeley.edu/pdfs/cdi_Brazil_report.pdf.

Eric Brewer is a professor of computer science at the University of California, Berkeley. He received a PhD in electrical engineering and computer science from MIT. Contact him at brewer@cs.berkeley.edu.

Michael Demmer is a PhD student at UC Berkeley, studying distributed systems and intermittent networking. He received a BS in computer science from Brown University. Contact him at demmer@cs.berkeley.edu.

Bowei Du is a PhD student at UC Berkeley studying distributed systems. He received a BS in computer science from Cornell University. Contact him at bowei@cs.berkeley.edu.

Melissa Ho is a research staff member at UC Berkeley. She received an MSc in data communications, networks, and distributed systems from University College London. Contact her at mrh9@cornell.edu.

Matthew Kam is a PhD student in computer science at UC Berkeley working on educational technology and human-computer interaction. He received a BA in economics and a BS in electrical engineering and computer science, also from UC Berkeley. Contact him at mattkam@cs.berkeley.edu.

Sergiu Nedevschi is a PhD student in computer science at UC Berkeley, studying ICT for developing regions and wireless networking. He received a BS in computer science from the Technical University of Cluj-Napoca, Romania. Contact him at sergiu@cs.berkeley.edu.

Joyojeet Pal is a PhD student in city and regional planning at UC Berkeley, where he received an MS in information management and systems. Contact him at joyojeet@cs.berkeley.edu.

Rabin Patra is a PhD student in computer science at UC Berkeley, studying architecture and wireless networking. He received a B.Tech (bachelor of

technology) in computer science and engineering from the Indian Institute of Technology, Kharagpur. Contact him at rkpatra@cs.berkeley.edu.

Sonesh Surana is a PhD student in computer science at UC Berkeley, studying networking and ICT for healthcare. He received a BS in computer science from Carnegie Mellon University. Contact him at sonesh@cs.berkeley.edu.

Kevin Fall is a senior networking researcher at Intel's Berkeley Research Laboratory. He received a PhD in computer science and engineering from the University of California, San Diego. Contact him at kfall@cs.berkeley.edu.



Innovative Technology for Computer Professionals Computer

Welcomes Your Contribution

**Computer
magazine
looks
ahead
to
future
technologies**

- **Computer**, the flagship publication of the IEEE Computer Society, publishes peer-reviewed technical content that covers all aspects of computer science, computer engineering, technology, and applications.
- Articles selected for publication in **Computer** are edited to enhance readability for the nearly 100,000 computing professionals who receive this monthly magazine.
- Readers depend on **Computer** to provide current, unbiased, thoroughly researched information on the newest directions in computing technology.

**To submit a manuscript for peer review, see
Computer's author guidelines:**

www.computer.org/computer/author.htm

