

Networking Technologies in the Developing Nations

Rick McGeer, Marc E. Fiuczynski, Jack Brassil, Larry Peterson, David Reed
Hewlett Packards Labs & Princeton University & MIT

A significant opportunity for the networking community is the deployment of information and communication technologies (ICT) in the developing nations. This is enabled by the coming convergence of long-range high-speed wireless technologies, which permit the deployment of widespread ICT without the requirement of laying wire for the expensive “last mile”, and the widespread availability of cheap, rugged end-user devices within the means of people in the developing nations and well-adjusted to variable environmental conditions. This development promises great and positive social impact while offering not simply the opportunity but the requirement to rethink the basics of network services and protocols.

The social impact of ICT in the developing nations is difficult to underestimate. A celebrated example is the deployment of cell phones in the villages of Bangladesh financed by microcredit from the 2006 Nobel Peace Laureate Muhammad Yunus and the Grameen Bank. The village “phone ladies” created profitable businesses by selling minutes to their neighbors, thus enabling communication with relatives overseas and in particular the arrangement of remittances. Over 50,000 villages have connectivity due to this project. A second example is the deployment of long-range 802.11 in southern India by UC Berkeley’s Tier research group. Using these links over 2000 patients/month receive eye care that was previously unavailable. About 10% of the patients achieve *significant* vision improvement, and 85% of those return to income generation. Current growth plans should cover a population of 2.5M in southern India alone.

The availability of both inexpensive wireless technologies and devices offer the opportunity for the citizens of the developing nations to jump several generations of ICT, going directly from analog wired circuit-switched connectivity (or none at all) to digital wireless packet-switched networks. However, these networks will be different from their counterparts in the industrialized nations. In particular, the following properties will hold:

- Last-mile connectivity will generally be provided by some form of wireless technology, typically mesh networking, and often with pure peer-to-peer networks. Current networking models focus on carrier-provided infrastructure with centralized service support authority. Access in developing nations may focus more on local connectivity provided by collaborative infrastructure. The questions to be answered are how one manages and troubleshoots a network where there is no single controlling and responsible authority. In this sense, the vision is close to that of the original Internet, as a cooperative medium among end hosts. In this case, however, we can assume that the peering gateways/routers are mutually untrusting.
- Connectivity in the developed nations, and in particular in the continental United States, assumes low latency links from the customer premise to the rest of the Internet. Connectivity in the developing nations may be episodic, lossy, and/or have high latency. Overlays providing delay-tolerant, content-based, and multi-path routing and transport are likely to be requisite features of such networking. In other words, intelligence and storage within the core of these networks will likely be a must—quite unlike the original Internet.
- Networking in the developing nations—and to some extent the developed nations—is becoming focused on applications that serve the person working in his or her social context. An excellent example use of ICT for efficient local market prices, which can increase rural income and reduce lost harvest (by picking at the right time). Thus, the goal may not primarily be providing access to the public Internet. An architecture that derives from such considerations is likely to be a set of peer-to-peer standards that permit collaboration over the wide area among social networks, while preserving end-user control over security, privacy, and collaboration.

Research into standards, protocols, and services with these properties is already taking place on PlanetLab. An excellent example is the long-running service OpenDHT, developed by Sean Rhea and his collaborators. Jacob and Scott Shenker demonstrated the value of OpenDHT as a collaborative medium with their EZShare.org photo-sharing service. In this service, OpenDHT was used as an indexing and discovery mechanism for a peer-to-peer photo-sharing service. In many ways, we believe that the EZShare.org service represents a prototype for a new class of applications. It is a service primarily offered by a community to itself; the service, after all, is the photographs, not the indexing service. Its basic backbone is an object sharing and transfer service, which can be applied to many sorts of objects, not simply photographs. It could, for example, form the foundation of a peer-to-peer electronic mail service. Finally, it demonstrates the ability of a wide-area platform to offer substantial resources to those who cannot command such resources by themselves; e.g., Jacob Shenker—a middle-school student—developed and deployed EZShare.org.

Another excellent example of peer-to-peer collaborative services is the OurGrid project in Brazil (<http://www.ourgrid.org>). OurGrid was conceived as a conventional Grid resource-sharing project for intense computational tasks; it has recently spread to lighter-weight, less batch-oriented applications. A specific service of significant interest is the OurBackup application, a peer-to-peer grid-based backup service.

Services such as EZShare.org and OurBackup are of course of interest to the developed nations as well as the developing nations. However, in the developed nations server-based commercial services often perform these tasks. The challenges for heavy-infrastructure commercial services in deploying these applications are now well understood by these providers. A light-infrastructure peer-to-peer approach is a far more challenging research area, and is likely to be critically important for networking in the developing nations – and to give us new insights into networking in the developed nations.

Case Study

In this section, we look at a proposed developing nation infrastructure and networking project as a case study, and examine the networking challenges it entails. The study is the One Laptop Per Child (OLPC) project, championed by Nicholas Negroponte of MIT. This project envisions massive networking deployments in challenging environmental conditions with primitive on-site support and management. In this, it represents a significant break from classic developed world networks, which rely on pristine environmental conditions in carefully engineered privately owned data centers and NOCs. In the developing world, in many cases, the endpoints are the infrastructure.

OLPC

The One Laptop per Child (OLPC) project is a strong case study for Real Overlays and Distributed Systems (ROADS) that can be explored on the existing PlanetLab facility available to distributed systems researchers and/or other facilities that more closely resemble an OLPC deployment.

Former MIT Media Lab Director Nicholas Negroponte conceived the One Laptop Per Child project. During a sabbatical in rural Cambodia, he gave each schoolchild in the local village an IBM laptop computer. Negroponte and his son Dmitri then installed a satellite Internet base station and a wireless access point in the village school. The laptops were tremendously popular in the village, not least because the screens were the best source of light in most village huts during the evening. They were also effective. The first English word the children in the village learned was “Google”. Moreover, the children showed dramatic educational progress by standardized test methods that Negroponte used to measure progress.

The dramatic success of the Cambodia village project persuaded Negroponte and his collaborators (Seymour Papert, Mitchell Resnick, and Alan C. Kay) to design a small, rugged, inexpensive, wireless laptop computer that governments could deploy to schoolchildren in lieu of textbooks. At a target price point of \$100 and a three-year lifetime, the OLPC laptop is billed as a cost-effective replacement for textbooks with, of course, the benefits of programmability and access to the greatest library the world has ever known – the Internet, with the accumulated knowledge of mankind a click and a few seconds away.

If successful, it is difficult to imagine another single achievable technological innovation that could achieve such a significant transformation in so many people's lives, so very quickly.

OLPC, if successful, would clearly be of great social benefit. At a stroke, it is a promise to bridge the digital divide and offer heretofore unattainable educational opportunities to the world's poor; an information infrastructure equal to that of the developed nations at a fraction of the cost; and, a foundational, rapidly deployable, robust information infrastructure technology for use in disaster-response situations.

The sheer scale of the OLPC effort is daunting. The typical deployment the project envisions is 50-500 laptops in a village or school, which form an ad-hoc wireless mesh network. The laptops are then connected to a gateway, which is connected by a backhaul link to the national education network (or commercial ISPs) and thence to the Internet. The *trial* deployment in most nations is expected to be 1,000,000 laptops with an expected distribution of roughly 10,000 sites/nation. The total run rate of manufacture is expected to be on the order of 165,000,000 units/annum, giving a total worldwide deployment of 500,000,000-1,000,000,000. These devices will, once deployed, form the single largest class of devices (with reasonable screen size) connected to the Internet.

The networking challenges faced by the OLPC project dwarf anything that has ever been attempted. While a conventional approach of attaching the OLPC laptops to the existing Internet might work, it is likely that meeting the networking challenge with existing solutions will result in a lesser Internet experience. We detail the challenges here:

- **The Village Gateways form the largest self-managing network ever conceived.** The OLPC project envisions a ratio of 100 laptops per gateway, with a total deployment ranging from 1000 gateways in a small nation such as Libya to almost 200,000 in Brazil. Each gateway will be deployed in an environment of challenging conditions (temperature ranges from -5°C to 45°C, ambient dust, humidity, and irregular and dirty power supplies), with minimally trained on-site support (typically, a teacher with little or no prior exposure to ICT). To keep total cost of ownership low, a small staff at one or more national network operations centers must have the ability to both remotely manage and remotely debug: the machines (i.e., gateways and laptops), the services running on those machines, the networks that connect them, and data they send/store. Today's PlanetLab deployment, with a machine count approaching 1,000 machines worldwide, is run by two operational staff members (assisted by another four R&D team members). Planetlab is the most ambitious, open source based, distributed autonomous-management enterprise undertaken in information technology: PlanetLab approximates *trial* deployments for the OLPC gateways.
- **The Village Mesh Networks will be the largest ever undertaken.** No one has formed an operational (as opposed to experimental) mobile ad-hoc networks of the scale conceived by OLPC (50-500 machines), particularly in the presence of anticipated significant environmental challenges. Management of such a network—particularly autonomic management from a village gateway—is a challenge that has never been undertaken.
- **The OLPC Content Distribution Network is the largest ever undertaken.** At its peak, the OLPC network envisions simultaneous content distribution (national textbooks, software, web content) to several hundred million devices through several million gateways. Today's largest content distribution network is two orders of magnitude smaller than that; the largest *public* content distribution network, CoDeeN, is four orders of magnitude smaller than that.
- **The OLPC project envisions the largest shared store ever conceived.** The OLPC devices, for reasons of cost, power consumption, and durability, do not contain a hard disk: 1-4GB of flash is the on-device store. Because the laptops might be lost or damaged, they must be backed up into the network. In many ways, the OLPC device is a cache for the network and the network a cache for the device. The village gateway is an inappropriate device for permanent storage: its hard drive and unattended nature combined with the environmental challenges leave it little less vulnerable than the personal device. We expect that techniques to persist data will be inclined to use peer-to-peer techniques as explored by OurBackup, the Logistical Computer Infrastructure (LoCI/LoDN) or OceanStore, which will likely be supplemented by a large national backing store.

- **The OLPC Project will do all this over the most heterogeneous network ever built.** Within rural areas, connectivity will generally be through satellite or cellular wireless links, having asymmetric bandwidth and high-latency connectivity: on the order of 250 milliseconds to 2.5 seconds RTT. Such links are also subject to significant outages due to environmental conditions, such as rain fade in the commonly-used Ku band or overloaded cellular towers. Satellite links may or may not be augmented with slow landline or long-haul wireless. In all cases, delay-tolerant networking is a priority; in cases where backup paths are available, fast switching (beyond what is achievable using conventional routing protocols) from failing or failed paths to working paths is required.

Summary

If we examine the list of challenges discussed here, the alignment with the GENI research agenda is apparent. We have the following research challenges that derive immediately from the various networking projects in developing nations for the GENI research community:

- autonomic management of large-scale distributed infrastructures,
- autonomic management of large mobile ad-hoc networks,
- provisioning and update of large scale distributed storage systems,
- large scale content distribution systems, and
- sensing and control of routing and transport over large, heterogeneous networks with high latency and link loss.

Networking projects in the developing world challenge both our technical capabilities and our social conscience. Collectively, they form a clarion call to our profession, our nation, and our civilization. Our choices are to meet this challenge and reinvent the Internet to rejuvenate the Enlightenment for another generation and expand it to those to whom the dark ages beckon, or to stand bereft before the bar of history.