ROADS for Developing Nations

Rick McGeer[1], Marc E. Fiuczynski[2], Jack Brassil[1], Larry Peterson[2], David Reed[1,3]
[1] Hewlett Packards Labs
[2] Princeton University
[3] MIT

This position paper is a call to arms to conduct research, development and deployment of ROADS (Real Overlay and Distributed Systems) for developing nations. The deployment of information and communication technologies (ICT) in the developing nations represents a significant opportunity for the distributed systems and networking community. This opportunity is enabled by the 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 significant social impact and thereby offers a great opportunity to rethink the basics of network services and protocols [BRE05].

The availability of both inexpensive wireless technologies and devices may let citizens of developing nations 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 (or should) likely be different from their counterparts in the industrialized nations. In particular, the following properties will hold:

- Last-mile connectivity will not be based on wireline broadband technology, but will generally be provided by some form of wireless technology, possibly mesh networking, and often with collaborative peer-to-peer networks.
- Connectivity in the developing nations may be episodic, lossy, and/or have high latency. Applications and services based on delay-resilient, 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.

As an example of a transformative vision of technology for the developing world, we look at the One Laptop Per Child (OLPC) project, championed by Nicholas Negroponte of MIT. This project envisions massive network deployments in challenging environmental conditions with novice on-site support and management. In this respect, it represents a dramatic 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.

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 provides an uplink to the national education network (or commercial ISPs) and thence to the Internet. The *inital* deployment in most nations is expected to be 1,000,000 laptops with an expected distribution of roughly 10,000 sites/nation. The global run rate of manufacture is expected to be on the order of 165,000,000 units/annum, giving a 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 distributed systems and networking challenges faced by the OLPC project dwarf anything that has ever been attempted. We detail the challenges here:

- **The Village Gateways form the largest self-managing network ever conceived**. The OLPC project estimates 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 -5ºC to 45ºC, 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). At this scale, PlanetLab approximates a *small* deployment of 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 WiFi 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 commercial content distribution network is two orders of magnitude smaller than that; the largest *research* content distribution network solutions deployed on PlanetLab, Princeton's CoDeeN and NYU's Coral, 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 research projects such as OurBackup, the Logistical Computer Infrastucture (LoCI/LoDN) or OceanStore, which may 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 both variable latency and long delay 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, solutions that are resilient to long delay and loss are a priority; in cases where backup paths are available, multi-path transport or fast switching (beyond what is achievable using conventional routing protocols) from failing or failed paths to working paths is required.

If we examine the list of challenges discussed here, the alignment with the agenda of future networking research (such as GENI or FIRE) is apparent. We have the following research challenges that derive immediately from the various projects in developing nations for the distributed systems and networking 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 nations challenge both our technical capabilities and our social conscience. Collectively, they form a clarion call to our profession, our respective nations, and our civilization. In the full length version of our paper we would like to present in further detail the networking properties that differ in developing nations vs. developed nations, map ongoing research on PlanetLab to the shorter-term challenges faced by developing nations making the leap with OLPC (or alternatives such as Intel's Classmate), and conclude with identifying in further depth the long-term research challenges.

References

[BRE05] Brewer, E., Demmer, M., Du, B., Fall, K., Ho, M., Kam, M., Nedevschi, S., Pal, J., Patra, R. and Surana, S. "The Case for Technology for Developing Regions." IEEE Computer, 38 (6). 25–38. 2005