TinyPC: Enabling Low-Cost Internet Access in Developing Regions

Muneeb Ali and Koen Langendoen
Delft University of Technology
Mekelweg 4, 2628 CD Delft
The Netherlands
{m.ali, k.g.langendoen}@tudelft.nl

ABSTRACT

Benefits of the Internet and communication technologies are limited to the world's "first-world" population. Cost factors, low literacy, and limited access to power and bandwidth in developing regions prohibit Internet access for potential developing-world users. Apart from limited communication infrastructure (telephone lines, ISPs, 802.11 long-distance links, etc.) a fundamental road-block in providing Internet access to developing-world users is the lack of affordable end-user computing devices. Most previous attempts to develop inexpensive computing devices were commercial failures as they were unable to keep the retail costs low. In this paper we present the design of TinyPC - a low-cost computing device specifically aimed at providing basic Internet access (web browsing, email) to developing-world users. TinyPC is inspired by recent advances in embedded networked systems (like sensor networks) and we show that, even today, with TinyPC it is possible to connect a developing-world user to the Internet in price ranges well below $100.

Categories and Subject Descriptors

J.7 [**Computers In Other Systems**]: Consumer Products

General Terms

Documentation, Design, Economics, Human Factors.

Keywords

Developing Regions, Low-cost Computers, Internet Access

1. INTRODUCTION

The Internet with its world-wide broadcasting capability, easy mechanisms for information dissemination and interaction, has revolutionized the way people communicate with each other, and the way people access information. However, these benefits are basically limited to the world's "first-world" population, e.g. according to a 2007 survey [1] Internet penetration in North America is 69.7% of population compared to 3.6% for Africa and 10.7% for Asia (Figure 1).

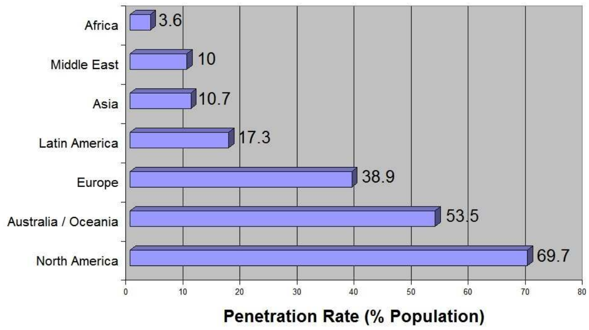

Figure 1: Internet penetration by world region [1]

In many developing-world countries, there is limited access to information as libraries are almost non-existent and there is no easy access to books. In places that have no other information resources, the Internet has the potential to deliver information electronically. However, challenges like cost factors, low literacy, intermittent access to power, narrow bandwidth (which translates to higher ISP and telephone charges), lack of reliable communication infrastructure, and limited access to computing resources prohibit Internet access for potential developing-world users. Although there are a number of unique technical, environmental, and cultural challenges in developing regions (see [11] for details), one of the *fundamental* road-blocks in providing Internet access to developing-world users is the lack of affordable end-user computing devices. Internet access in developing regions (over dialup) is relatively cheap (e.g. a $1 Internet "scratch" card would last for 10 hours) but purchasing a computer is a large initial investment.

Brewer et al. [10] argue that even though technology can play a large role in developing regions, there is a need for technology research specifically aimed at developing regions as technologies developed for the "first-world" have been a poor fit in these areas. This phenomena is also observed in the slow adoption of computing devices (desktop computers, laptops, PDAs) by developing-world users. On the other hand, attempts at developing inexpensive computing devices specially for developing regions (Simputer, Computador Popular) were also commercial failures [18]. Initiatives like the much-hyped $100 laptop and Intel's program to develop a cheap computer are yet to have any real impact.

NSDR'07, August 27, 2007, Kyoto, Japan.

In this paper we present our vision (Section 2) of a low-cost computing device, called TinyPC, which is designed to provide basic Internet access (web browsing, email) to developing-world users. The basic idea behind TinyPC is to use very small inexpensive devices for computing (inspired from advancements in networked embedded systems like sensor networks [13]) and use television (available in many developing-world houses) for display. We discuss the feasibility of our vision in Section 3 and show that, even today (2007), it is possible to connect a developing world user to the Internet in price ranges well below $100. In the coming years the number of small embedded devices connected to the Internet would outnumber the traditional Internet hosts [12] and given such large production scales of tiny networked devices the price of such devices could even drop well below $50 in the future.

We outline design goals in Section 4, present the design of TinyPC in Section 5, and discuss open issues and future research challenges in Section 6

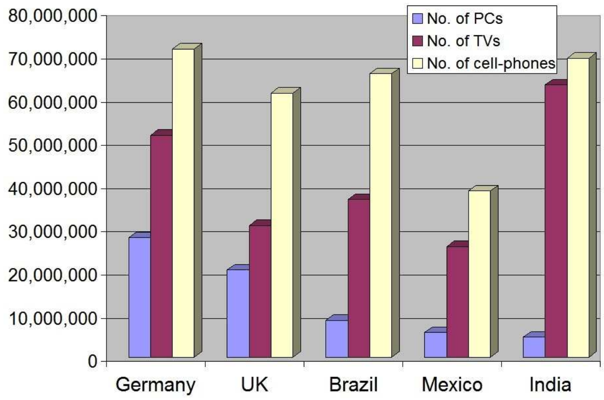

Figure 2: Total number of PCs, TVs, and cell-phones [2]

2. VISION

To date Moore's law, the guiding principle that has driven the computing industry, has proven remarkably accurate. Computer chips continue to become faster, smaller, and cheaper. This is already changing the way people access the Internet; you do not need a traditional desktop computer to go online, even a small device like a cell-phone can provide Internet access. The sheer number of cell-phones sold every year and the ever decreasing costs of handsets seem to imply that in the near future, cell-phones may become the primary computing device for Internet access [21].

Future cell-phones will have all capabilities of a traditional desktop only with two short comings - the display and keyboard are small [6]. In theory if this cell-phone of the future is connected to an external display and a keyboard the entire user experience of a traditional desktop can be reproduced. Such powerful future cell-phones with massive USB storage, can replace PDAs, laptops, and desktops as they are easier to carry than laptops and with external keyboard and display they can function similar to desktop computers [6].

At first glance, this vision of cell-phones as primary computing devices seems to have important implications for developing-world countries. Cell-phones and televisions, unlike computers, have been a great success in developing regions (e.g. India, Pakistan) and the number of people with cell-phones and televisions is growing at a faster rate than users with computer access. Figure 2 shows the total number of computers, televisions, and cell-phones in some developed (Germany, UK) and some developing (Brazil, Mexico, India) countries. The statistics about TVs and cellphones are from the CIA World Factbook, and the estimates for PCs are made by ITU (data available online at [2]).

The countries in Figure 2 are ordered (from left-to-right) by decreasing number of PCs and we observe that there is a sharp contrast between access to computers and number of TVs and cell-phones in developing regions, e.g. India. If a combination of cell-phone and TV can enable Internet access then we would see a significant increase in the number of Internet users in the developing-world. However, although we believe that "fancy cell-phones" of the future (as described in [6]) may become the primary computing device for Internet access in the developed countries, the high costs (of such powerful cell-phones and Internet access over cellular lines) will prohibit the adoption of this model in developing countries.

People in developing countries do have cell-phones, but they are mostly low-cost handsets that provide the bare-minimum (voice and SMS) functionality. Cell-phone adoption was a success in such areas because of the availability of *affordable* handsets. A cell phone is essentially a battery-powered microprocessor, providing a general-purpose computing platform, and wireless interface optimized for voice I/O. Manufacturers are able to keep the costs low for the bare-bones handsets (popular in developing regions) by leaving out sophisticated features like camera, local storage, full-color screen, and support for multiple wireless interfaces (e.g. bluetooth). These handsets focus on the core use of cell-phones and leave out other functionalities (including Internet support).

Building upon the lessons learnt from the success of cell-phones in developing regions, we argue that computers are most useful in developing regions as Internet access platforms and other uses of computers are of less importance. Once users have bare-bones Internet access (web browser and email) they can use free online resources like calculators, online storage (on web-based mail accounts), spread sheets and word processing tools (e.g. by Google) to satisfy most of their computing needs. The core functionality of such bare-bones computing platforms for developing regions should be to enable information access to the Internet and to provide access to cheap communication facilities like email.

The most important factor in enabling Internet access for developing-world users is the *cost*. Laptops and PCs are designed to do much more than just provide Internet access and because of their high price to-date they have been unsuccessful in bringing potential developing-world Internet users online. Now, imagine a cheap bare-bones computing device that is optimized for web browsing and email much like low-cost cell-phones are optimized for voice I/O and SMS - we call that device TinyPC. Users will connect TinyPC to a television (available in many developing-world homes as shown in Figure 2) for display, leverage the internal wireless (802.11) connectivity of TinyPC for Internet access, and use cheap locally produced mouse and keyboard as input devices.

3. FEASIBILITY

Before we present the feasibility of our TinyPC vision, we briefly discuss the characteristics of target developing-world environments.

3.1 Target Environment

TinyPC is mainly targeted towards the urban population of developing world countries i.e. people who have (or can easily gain) know-how to operate computers and surf the Internet, but simply do not have enough money to purchase their own computers. As argued earlier (Section 2) television penetration in developing countries is quite significant, and if a user has access to a television display then he also has access to (in the worst-case intermittent) electricity. The resources missing for Internet access, in such environments, are cheap computing platforms and some form of communication medium (although telephone lines for dialup or DSL connections are largely available).

The challenges in rural areas of developing countries are much greater. Telephone connectivity suddenly drops to almost non-existent as compared to urban areas, literacy rate also drops significantly and we cannot assume that people in villages can use input devices or surf the Internet. Our focus is not on users who (due to literacy problems) cannot benefit from the Internet but rather on people who can (and want to) go online but monetary factors limit their options.

3.2 Technology Trends

TinyPC is inspired by recent hardware and software advancements and in this section we show how by utilizing these advancements (in a single product) developing world users can go online with extremely low-costs.

3.2.1 Hardware

A given computing capacity becomes exponentially smaller and cheaper with each passing year, following Moore's law. Recently, this lead to a new generation of cheap embedded networked platforms like sensor networks (sensornets) [13]. A typical sensor node consists of sensing hardware, memory, battery, embedded processor, and transceiver and the hardware trends in sensornets show a steady progress from 8-bit (Mica [19]) and 16-bit (Telos [25]) platforms to ones with 32-bit processors (Sun SPOT [5]). The wireless radios on these sensor nodes show a trend from low-capability CC1000 radios to more sophisticated 802.15.4 radios [7]. Even though sensornets were designed with a completely different purpose in mind (sensing and reporting), researchers have used these devices to run TCP/IP and connect them to the Internet (see Section 3.2.2). These sensor nodes (with capabilities to connect to the Internet) are considered to be cheap-enough that they are discarded once the battery runs out.

Furthermore, recent research in 802.11 long-distance wireless networking [8, 24] has paved the way for affordable wireless connectivity in developing regions and projects like "Radio Free Intel" are aiming to integrate low-cost radio capability into every silicon product and making wireless connectivity ubiquitous for every device (no matter how small or cheap) that has a processor.

3.2.2 Software

Dunkels developed a micro-TCP/IP stack (uIP) for 8-bit architectures [14] (a full implementation of TCP/IP) and tiny sensor nodes run uIP to directly communicate with IP networks. Also, there are GPL implementations of IETF 6lowpan standard (e.g. by Sensinode a startup company in Europe) that provide IPv6 connectivity to tiny sensornet nodes and make them "first-class" IPv6 citizens of the Internet.

The Contiki operating system [15] (build around the uIP stack) is a small operating system that has been ported to various sensornet platforms and other devices with limited resources like Commodore 128, Game boy, Atari, and Atmel 8-bit AVR. The Contiki web browser (using uIP) is probably the first web browser ever to run on an over 25 years old computer system, e.g. 1979 Atari 800. If it is possible with with current softwares to connect such old and resource limited machines to the Internet then it is easy to image how a more powerful product like TinyPC can easily be used as a computing device for Internet connectivity.

Furthermore, research in delay tolerant networking [16] has enabled meaningful connectivity even under the presence of network failures and disturbances, and browsers like TEK [28] enable users to search the web off-line using email if Internet connectivity or cost prohibits online web search.

4. DESIGN GOALS

In this section we present the design goals critical to the success of computing platforms developed specifically for under-privileged environments:

- **Low Cost:** Cost is the *single most important* design goal in developing regions and the price tag will be the make-or-break factor in the success of any computing platform tailored for these environments.
- **Leverage Local Resources:** There is a need to pay attention to unique social and environment conditions and make full use of any resources that are already available in developing regions.
- **Only Essential Features:** Extra features will not only confuse the largely illiterate population but would also drive the costs up. There is a need for simple, easy to use computing devices that focus only on the core functionality.
- **Internet Access:** The main purpose of computing devices for developing regions should be to enable information access to the Internet. Developing-world users could benefit the most when connected to the ever-evolving Internet and will find little use for standalone computing platforms.
- **Wireless Connectivity:** Telephone connectivity remains extremely low in rural areas of developing countries and inadequate telephone lines in urban areas present a classic "last-mile" problem to fast Internet access. These regions will skip the dialup and DSL stage of the Internet and jump directly to wireless (802.11 or WiMax) Internet access. Computing platforms for developing regions should be optimized for wireless connectivity.

5. TINYPC DESIGN

In this section we present the design for TinyPC, we discuss design tradeoffs, cost issues, and compare TinyPC with previous attempts at designing low-cost computing devices for developing regions.

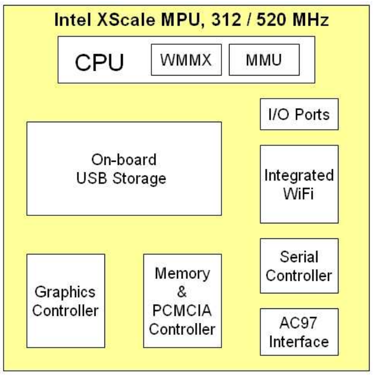

Figure 3: Simplified block-diagram of TinyPC

5.1 Design Tradeoffs

In Section 3.2.2 we argued that it is possible to connect tiny sensornet devices with 16-bit processors, limited memory, and low-power radios to the Internet. However, such sensornet platforms, e.g. Telos [25] are not the right tools for our purpose. They simply do not have the right facilities to generate an appropriate user interface. They can generate TV signals for standard NTSC/PAL TVs and text based displays, but being able to do VNC or any point-click will be difficult with a small 8-bit or 16-bit micro-controller. Therefore, we opt for a higher-powered (32-bit) processor - more specifically the XScale processor from Intel (although any ARM32 processor should be sufficient). The XScale processor is more capable than the processors commonly found on sensornet nodes, can do DSL/Ethernet, and has a display hardware accelerator. For wireless connectivity, instead of 802.15.4 radio used by sensornet nodes, we opt for 802.11 wireless radio as 802.15.4 radio is short range and has limited bandwidth.

5.2 TinyPC Hardware

TinyPC can be thought-of as a combination of a higher-power sensornet node like Sun SPOT [5] or Intel Imote2 (without any sensing hardware) and a cheap USB drive. Figure 3 shows a simplified block diagram of the TinyPC platform. Below we briefly describe different components of TinyPC:

5.2.1 Computing

As argued in Section 5.1 we choose Intel Xscale 32-bit processors for TinyPC. More specifically in the current design we choose the Intel XScale IXP420 network processor. Intel IXP4XX series processors focus on lower cost, scalable performance, and reduced power and is suitable for residential and small/medium enterprise network applications.

5.2.2 Storage & Memory

Processors, radios, and other computer components are not improving greatly in cost, but non-volatile memory is. The ability to cheaply store large amounts of data locally on TinyPC will greatly help the potential users as they move between home and workplace/school. The current design of TinyPC includes 1 GB of USB-based storage (integrated on board) and essentially TinyPC could be used as a USB drive as well. Apart from the USB storage, the current design of TinyPC has 16 MB DRAM and 2 MB Boot FLASH.

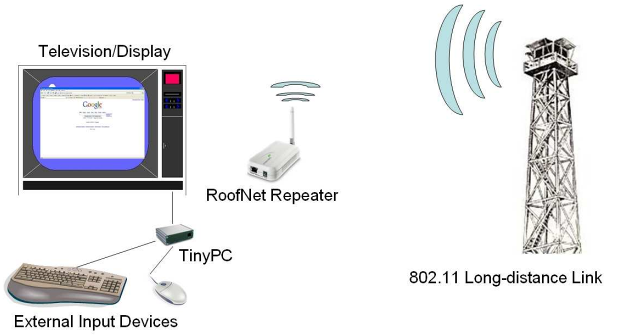

Figure 4: Wireless Internet access with TinyPC

5.2.3 I/O

The goal of TinyPC is to provide "desktop-like" Internet experience to users. Therefore, TinyPC is designed to have external I/O devices - keyboard and mouse for input and a television for display. Our goal is to leverage local resources in developing regions as much as possible and (as shown in Section 2 and Figure 2) TVs are available to a large population of developing countries. Utilizing the already available TVs greatly helps reduce the cost of TinyPC and the research challenges of using TVs as display devices are discussed in Section 6.

5.2.4 Networking

As argued in Section 3.2, 802.11 long-distance wireless networking [8, 24] brings connectivity in developing-world areas that have no other communication infrastructure. For the networking component we choose to integrate 802.11b/g compatible radio on the TinyPC board. Figure 4 shows how a typical user would connect to the Internet using TinyPC, a TV for display, and mesh-networking protocols like MIT RoofNet [9] (commercialized through Meraki [3]) that connect to 802.11 long-distance links (in case of rural access) or to some available ISP (in case of urban access).

5.3 TinyPC Software

The software trends inspiring the design of TinyPC like Contiki OS [15] and uIP stack [14] (described in Section 3.2.2) are open source softwares. We do not opt for using Contiki, which is optimized for running on limited resources, because with the Xscale processor and considerable RAM and storage it is easily possible to support Linux instead. TinyPC is envisioned to be a completely open source software platform. Free GPL software make the most sense in developing regions as software costs add up to a large portion of the total cost of owning a computer. Further, encouraging open source software will help fight software piracy problems common in developing regions.

5.4 Cost

Our aim is to produce TinyPC with a price tag of $50. The XScale processor costs about $20 for the processor. 1 GB stand-alone USB drives are available for $19 today (2007). For TinyPC we integrate the USB on board and this helps reduce the production and packaging costs, i.e. we hypothesize that the on-board integrated USB-storage of TinyPC will cost much less than $19. Keyboard, mouse and cables are significantly cheaper as, in most cases, they are pro-

Name	Year	Memory	Processor	Storage	Network	Display	Price Target/Actual
Computador Popular	2001	64 MB RAM	AMD K6-II 500 MHz	Flash 32MB	Ethernet, 56K Modem	Monitor	$300 (Target)
Simputer	2006	32 MB DRAM	Intel StrongARM	Flash 32MB, upto 128MB	Internal Modem, IR, and WLL	Build-in LCD	$275 - $440 (Actual)
MIT $100 Laptop	2006	256 MB DRAM	AMD LX-700, 433 MHz	Flash 1024 MB	802.11b/g	Build-in TFT	$100 (Target)
TinyPC	2007	16 MB DRAM	Intel Xscale, 312-520 MHz	2 MB FLASH, 1 GB USB	802.11b/g	TV or External Display	$50 (Target)

Table 1: Comparison of developing-world computing platforms

duced locally in developing countries, e.g. in Pakistan it is possible to purchase a keyboard and mouse for just $5. It is important to keep in mind that most prices quoted here are bill of materials ("BOM") prices and the end retail product generally costs 20-25% more than the total price of all the individual components.

To hypothesize the price of TinyPC today (2007) we look at a somewhat similar product - Moteiv's Telos Motes [25]. Telos uses a MSP430 micro-controller (16-bit RISC processor) and a packet radio compliant with the IEEE 802.15.4 standard, and a single unit costs around $60 - $70. With a $20 XScale processor and support hardware, it is possible to produce TinyPC in the $70-80 range, even today, depending on volume. For Internet connectivity with TinyPC, wireless mesh-networking enablers like Meraki Mini [3] are available for $49 and their price will go down with growth in use. It is important to note that wireless mesh-networking repeaters are meant for shared instead of individual use.

5.5 Comparison

TinyPC bears resemblance to some old approaches. In the late 90s, Sun, IBM, and others built a series of network computers (NCs). The NCs had bare bones capabilities and booted from a networked server and this reduced the cost per employee at large organizations. The Atari 2600, Commodore 64, and others leveraged the fact that people had TVs in their home. Commodore is an example of using TVs as the display to a generic computing environment. They did this to keep the price low to open up computers to high volume consumer usage and to leverage a familiar environment. Low-cost gaming devices with 80's home-computer hardware and TV as display (Sega MegaDrive/Genesis, Atari, or Commodore 64) are sold for less than $50 in toy stores today. Contiki has shown that running Internet-applications on such gaming-console hardware is feasible. The low price of such gaming-consoles (well below $50) supports our hypothesis of a network-connected TinyPC for $50.

Most previous attempts to develop inexpensive computing devices, specifically for developing regions, were commercial failures due to varying reasons (see [17, 18] for details). Two popular examples of previous attempts at designing computing platforms for developing regions are the Brazilian Computador Popular and the Indian Simputer. Hardware specifications for the different developing-world computing platforms are given in Table 1.

The Computador Popular (CP) was supposed to be a $300 device that comes in two modes, stand-alone and networked, and they wanted to keep the price down by using off-the-shelf components, possibly of the previous generation and by removing the hard drive. As of 2003, it was never produced due to local Brazilian political factors beyond the control of the government. Compared to TinyPC the CP was a "conventional computer" without a hard drive and we argue that removing the hard disk was not a good design choice since prices of storage are falling much more rapidly than other computer components.

The case for the Indian Simputer is also similar only that Simputer was actually produced and came to market at a price of $240 to $440 (three times the price of a low-end PDA). We believe that in the case of the Simputer the price of the integrated PDA-like display is making it expensive for potential users. Also, keyboards and big displays are natural and preferred ways of interacting with computers and potential users show resistance to the Simputer because of its I/O interfaces (smaller display, touch-screen).

The much-hyped MIT $100 laptop [4] is primarily aimed at children and educational use. Unlike the CP there is much technology innovation behind the $100 laptop. $100 laptop wants to drive down the price of the display, TinyPC does not have a display at all. A display is a big power sink and it drives up the cost of the power-supply. By not including a display in TinyPC we can have a cheaper power-supply as well. We acknowledge the fact that not all users already have a TV display, so the cost of display is not zero. However, TinyPC is targeted towards individuals who already own a TV (which is a considerably large population in developing regions). Another problem with using TVs as display is that TVs are not easily portable; we discuss this issue in Section 6.

Compared to web access via cell phones, TinyPC will enable a closer to "desktop-like" Internet access experience with large colored displays (depending on the TV screen), proper keyboard and mouse (compared to limited keypad of cell phones), lower price of device (to use the Internet over a cell phone you need a "fancy phone"), lower price of Internet access (the cell-phone industry charges a significantly higher price per-byte-of-data than ISPs). Over time "dumb" phones will get smarter and "fancy phones" will get cheaper, but both these trends will require considerable time before the benefits can reach the masses. With TinyPC developing-world users can be connected to the Internet with low-costs even today.

The research literature promotes the *shared computing* model for developing regions and Internet kiosks are widely suggested [22, 26] as efficient means of enabling information access in developing regions. TinyPC, however, targets a different problem and is aimed at personal or family use instead of shared usage. We believe that a combination of Internet kiosks and TinyPC (depending on individual situations) can bring many developing-world users online.

6. DISCUSSION

TinyPC opens up a breath of interesting technical and non-technical issues. Social issues in developing regions can work for or against new technologies, e.g. there is considerable resistance to technology in developing regions and TinyPC can experience such resistance. On the up-side it is common practice in places like India and Pakistan to give brides new TVs as dowry. Another issue is that, by connecting the TV to the computer the user's learning endeavors directly conflict with his family's only entertainment option. However, the rapid growth of gaming consoles (that also share the TV) indicate that sharing of TV within close family members might not be a big hinderance.

Labor costs in developed and developing regions differ significantly and for TinyPC to be a success it would eventually need to be developed locally. Appropriate marketing models, e.g. selling to governments that can then distribute to students, will prove important in the success of TinyPC.

LCD prices are coming down, and a small 4x6 LCD can cost about $15. To include a small display in future generations of TinyPC is a debateable topic. An issue with the TV approach is that the community for which TinyPC is being targeted also has old TV sets (in some cases black and white). Essentially, there is a need for TinyPC to seamlessly inter-operate will all TV sets. Is the effort in developing TV output for TinyPC worth the price reduction of removing a display? TV output will be an interesting area of future research and options like wireless to TV can be considered.

Research in 802.11 long-distance networking [8, 24] becomes more important with TinyPC. Apart from basic Internet access, TinyPC can find applications in education [20], low-cost health care [27], and speech recognition [23] in developing regions - all of which remain open research areas.

7. CONCLUSIONS

TinyPC can bring potential developing-world users online at low-costs and enable information access in developing regions. TinyPC is a new computing product and not a short term solution. As the envisioned "fancy cell-phones" (with external I/O) for Internet access in the western world become cheaper and new features are added to TinyPC, the difference between the two will reduce over time and eventually the two products may even become one - but till that day TinyPC can serve as a practical low-cost Internet access computing platform in developing regions.

Acknowledgements

The authors would like to thank Adam Dunkels (SICS), Zeeshan Syed (M.I.T), and Joe Polastre (Moteiv) for helpful discussions that shaped some of the ideas presented in this paper. Special thanks to Jan Beutel (ETH Zurich), Xiaofan Jiang (UC Berkeley), Joe Polastre, and anonymous NSDR reviewers for helping us improve the camera ready version. This work was supported in part by Project "Relate" No. 013790 funded by the European Commission (FP6 IST Programme).

8. REFERENCES

[1] http://www.internetworldstats.com/.

[2] NationMaster Website at URL: http://www.nationmaster.com.

[3] Meraki Networks, Inc. Website at URL: http://meraki.net.

[4] MIT $100 Laptop Website at URL: http://laptop.org.

[5] Sun SPOT System: Turning vision into reality. White Paper, June 2005.

[6] J. J. Barton, S. Zhai, and S. B. Cousins. Mobile phones will become the primary personal computing devices. In *7th IEEE Workshop on Mobile Computing Systems & Applications (WMCSA'06)*, pages 3–9, Washington, DC, USA, 2006.

[7] J. Beutel. Metrics for sensor network platforms. In *ACM RealWSN'06*, Uppsala, Sweden, June 2006.

[8] P. Bhagwat, B. Raman, and D. Sanghi. Turning 802.11 inside-out. *ACM SIGCOMM Comput. Commun. Rev.*, 34(1):33–38, 2004.

[9] S. Biswas and R. Morris. ExOR: Opportunistic multi-hop routing for wireless networks. In *ACM SIGCOMM'05*, pages 133–144, Philadelphia, Pennsylvania, USA, 2005.

[10] E. Brewer, M. Demmer, B. Du, K. Fall, M. Ho, M. Kam, S. Nedevschi, J. Pal, R. Patra, and S. Surana. The case for technology for developing regions. *IEEE Computer*, 38(6):25–38, June 2005.

[11] E. Brewer, M. Demmer, M. Ho, R. J. Honicky, J. Pal, M. Plauche, and S. Surana. The challenges of technology research for developing regions. *IEEE Pervasive Computing*, 5(2):15–23, 2006.

[12] D. Clark, C. Partridge, R. Braden, B. Davie, S. Floyd, V. Jacobson, D. Katabi, G. Minshall, K. Ramakrishnan, T. Roscoe, I. Stoica, J.Wroclawski, and L. Zhang. Making the world (of communications) a different place. *ACM SIGCOMM Computer Communication Review*, July 2005.

[13] D. Culler, D. Estrin, and M. Srivastava. Guest Editors' Introduction: Overview of sensor networks. *Computer*, 37(8):41–49, 2004.

[14] A. Dunkels. Full TCP/IP for 8 bit architectures. In *Proc. MobiSys'03*, San Francisco, May 2003.

[15] A. Dunkels, B. Grönvall, and T. Voigt. Contiki - a lightweight and flexible operating system for tiny networked sensors. In *Proc. of the First IEEE Workshop on Embedded Networked Sensors*, Florida, USA, Nov. 2004.

[16] K. Fall. A delay-tolerant network architecture for challenged internets. In *ACM SIGCOMM'03*, pages 27–34, Karlsruhe, Germany, 2003.

[17] R. Fonseca and J. Pal. Bringing Devices to the Masses: A comparative study of the brazilian computador popular and the indian simputer. In *South Asia Conf., UC Berkeley: Trends in Computing for Human Development in India*, pages 459–468, UC Berkeley, February 2005.

[18] R. Fonseca and J. Pal. Computing Devices for All: Creating and selling the low-cost computer. In *1st International Conference on ICT and Development*, UC Berkeley, May 2006.

[19] J. Hill and D. Culler. Mica: A wireless platform for deeply embedded networks. *IEEE Micro*, 2002.

[20] M. Kam, D. Ramachandran, U. Sahni, and J. Canny. Designing Educational Technology for Developing Regions: Some preliminary hypotheses. In *3rd International Workshop on Technology for Education in Developing Countries (TEDC)*, July 2005.

[21] S. Keshav. Why cell phones will dominate the future internet. *ACM SIGCOMM Comput. Commun. Rev.*, 35(2):83–86, 2005.

[22] R. Kuriyan, I. Ray, and K. Toyama. Integrating Social Development and Financial Sustainability: The challenges of rural kiosks in kerala. In *1st International Conference on ICT and Development*, UC Berkeley, May 2006.

[23] S. Nedevschi, R. Patra, and E. Brewer. Hardware speech recognition on low-cost and low-power devices. In *41st Design Automation Conference*, June 2005.

[24] R. Patra, S. Nedevschi, S. Surana, A. Sheth, L. Subramanian, and E. Brewer. WiLDNet: Design and implementation of high performance wifi based long distance networks. In *ACM/USENIX NSDI'07*, pages 87–100, Cambridge, MA, USA, April 2007.

[25] J. Polastre, R. Szewczyk, and D. Culler. Telos: Enabling ultra-low power wireless research. In *SPOTS track of IPSN'05*, Los Angeles, CA, USA, April 2005.

[26] A. Seth, D. Kroeker, M. Zaharia, S. Guo, and S. Keshav. Low-cost communication for rural internet kiosks using mechanical backhaul. In *ACM MobiCom'06*, pages 334–345, Los Angeles, CA, USA, 2006.

[27] Z. H. Syed. MIT Automated auscultation system. Thesis (M. Eng.) MIT EECS, May 2003.

[28] W. Thies et al. Searching the world wide web in low-connectivity communities. In *11th Int. World Wide Web Conference Global Community Track*, May 2002.