

Pinging Africa¹

R. Les Cottrell, SLAC

Introduction

Africa is huge with an area greater than Argentina, Europe, China, India, and the United States combinedⁱ. It has over 1 billion people or one seventh of the world's population (second only to Asia), speaking over 1000 languages, with climates from huge rainforests to vast deserts.

However, as a whole it suffers from many challenges including disease, conflict, corruption, protectionist policies (for example current cable and satellite companies have a lot to lose and many of these have close links to regulators and governmentsⁱⁱ), poverty, lack of skills, and poor infrastructure such as lack of power, inadequate telecommunications, inadequate Internet access, lack of PCs etc.

Africa also badly lags many years behind the rest of the world in its Internet performance, and has been falling further behind each year. In 2010 all of Africa combined had one third the international capacity as Austria aloneⁱⁱⁱ. The cost of Internet access for universities in Africa is 100 times more expensive than for U.S. universities^{iv}. The cost for an African university to buy the same bandwidth as a U.S. university (1 Giga-bit/second would be \$13M/month). Less than 12% of the Africa population in 2011 had Internet access. This percent is about 7 times less than for North America^v. For most of Africa it is even worse since most of the African Internet users are concentrated in Mediterranean Africa and South Africa as is seen in the cartogram showing country areas re-sized by Internet users^{vi}.

Measuring Internet performance and PingER

The PingER project^{vii} was launched at the start of 1995 to make Internet performance measurements initially for the SLAC National Accelerator Lab's High Energy Physics worldwide collaboration. By 1999 PingER had been extended to other HEP collaborations and there were 20 monitoring hosts in North America, Europe and Japan, and target hosts at over 350 sites in 54 different countries. In 2002 and 2003, in collaboration with the International Centre for Theoretical Physics (ICTP) in Trieste, Italy, and encouraged by the International Committee for Future Accelerators (ICFA) the monitoring was extended to cover developing countries. The current deployment extends to over 75 monitoring hosts in 23 countries, monitoring targets at over 800 sites in 166 countries. These countries contain about 99% of the world's population. The PingER deployment is shown in the Figure below:

¹ Work supported in part by the U.S. Department of Energy under contract number DE-AC02-76SF00515.
SLAC National Accelerator Laboratory, Menlo Park, California, 94025

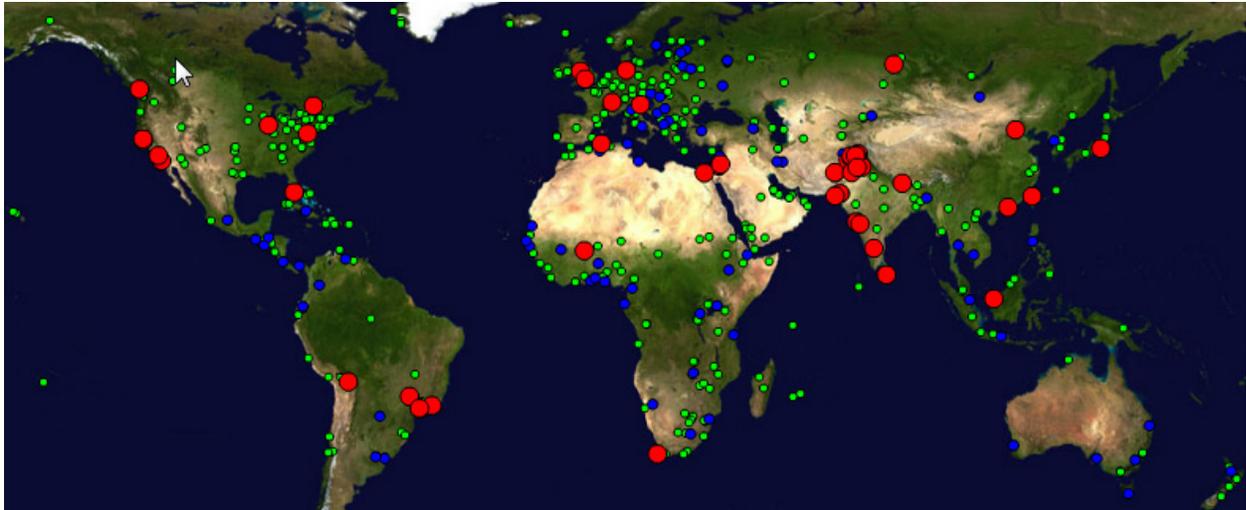


Figure 1: Locations of PingER monitoring and target sites as of June 2011. Red sites are monitoring sites, blue sites are targets at beacon sites that are monitored by most monitoring sites, and green sites are target sites that are monitored by one or more monitoring hosts

Each PingER monitor uses the ubiquitous Internet ‘ping’ facility to send a simple echo request packet to a target host. The target then simply echoes a response back to the monitor. From such a simple mechanism we can directly measure whether the target is responding at all (we call this the unreachability), the Round Trip Time (RTT) between sending the request and receiving the response, the percentage of request packets that are not responded to (the loss), and the variability of the RTTs (jitter). Each half hour each PingER monitor sends ten pings to each target in its list and records the results. The resulting measurements are stored locally by the monitor, and automatically gathered nightly, from each of the monitors, by three archive sites at SLAC at the Stanford University, the Fermi National Accelerator Lab (FNAL) in Chicago, and the National University of Sciences and Technology (NUST) in Islamabad, Pakistan. The online data goes back to the start of 1998, provides a fascinating history of Internet performance seen from multiple vantage points worldwide, and its volume is measured in TeraBytes.

Besides archiving the data, the archive sites also analyze the direct measurements to derive: the throughput; and the International Telecommunications Union’s (ITU) Mean Opinion Score (MOS) that provides a numerical indication of the perceived quality of received media such as Voice over IP (VoIP). The archive sites also aggregate the data into different time windows (hourly, daily, monthly and yearly) and groups such as hosts in a country, or in a region, or affinity groups (e.g. all hosts in countries belonging to the Arab League).

Worldwide internet performance seen by PinGER

Looking at the PingER information provides fascinating views of how the Internet is performing both in time and space. For example the packet losses tend to be mainly at the edges rather than within the backbones, and are thus relatively independent of the

distance between the monitor and the target. Losses of one or more percent can make real-time communications such a VoIP unusable. The losses seen from SLAC to regions of the world for the last 14 years are seen in the figure below.

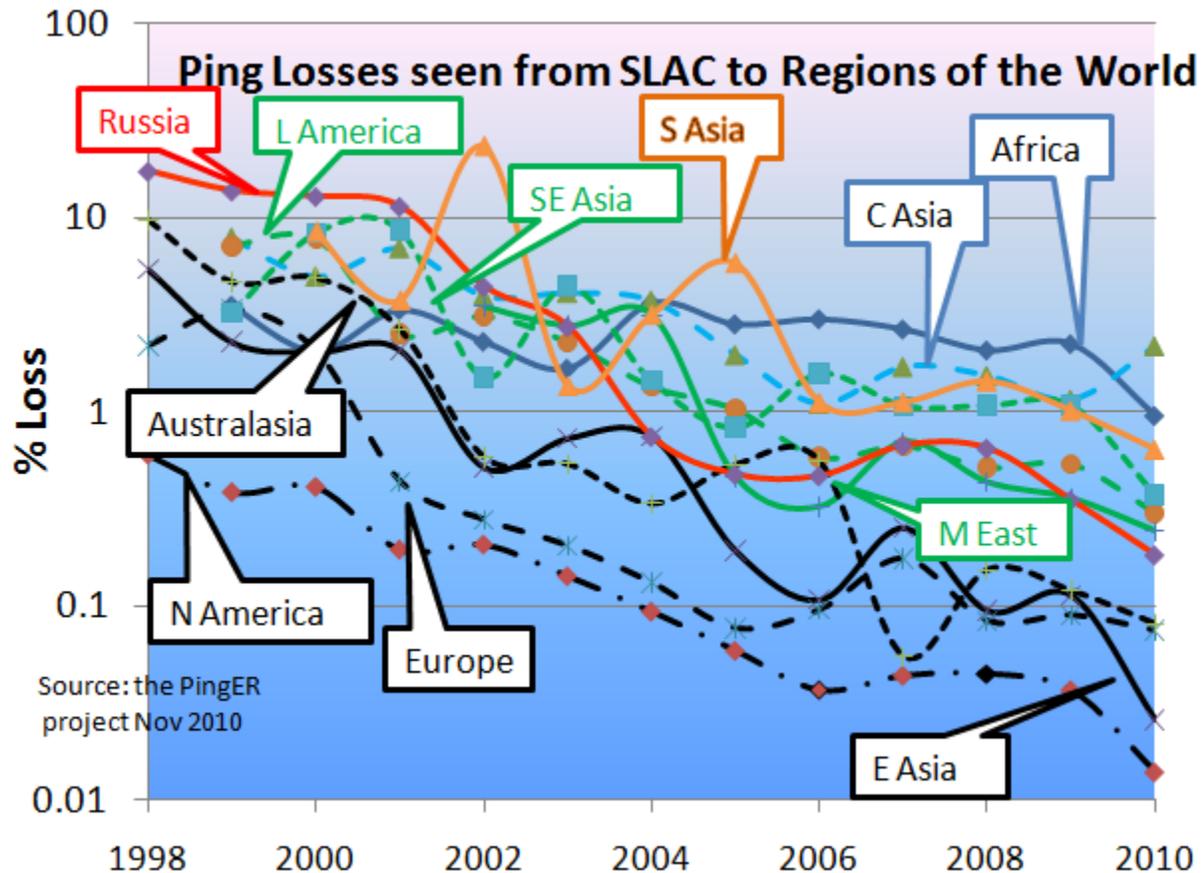


Figure 2: Packet Loss measured for various regions from SLAC in Nov 2010

The reduction in losses with time is roughly exponential, with an improvement of roughly a factor of 100 in 10 years. The world divides into at least two super-regions: the developed world (N. America, Europe, E. Asia and Australasia drawn in black), and the rest of the world. It is also seen that Africa and Central Asia in blue not only have the highest losses but their rate of improvement lags behind the rest of the world.

Similar results are seen for the derived throughputs. These are proportional to the inverse of the square root of the loss and the inverse of the RTT. Since the derived throughput is distance dependent we do not show the results for derived throughput from SLAC to North America in the figure below. Also rather than showing the individual data points we have show exponential trend lines.

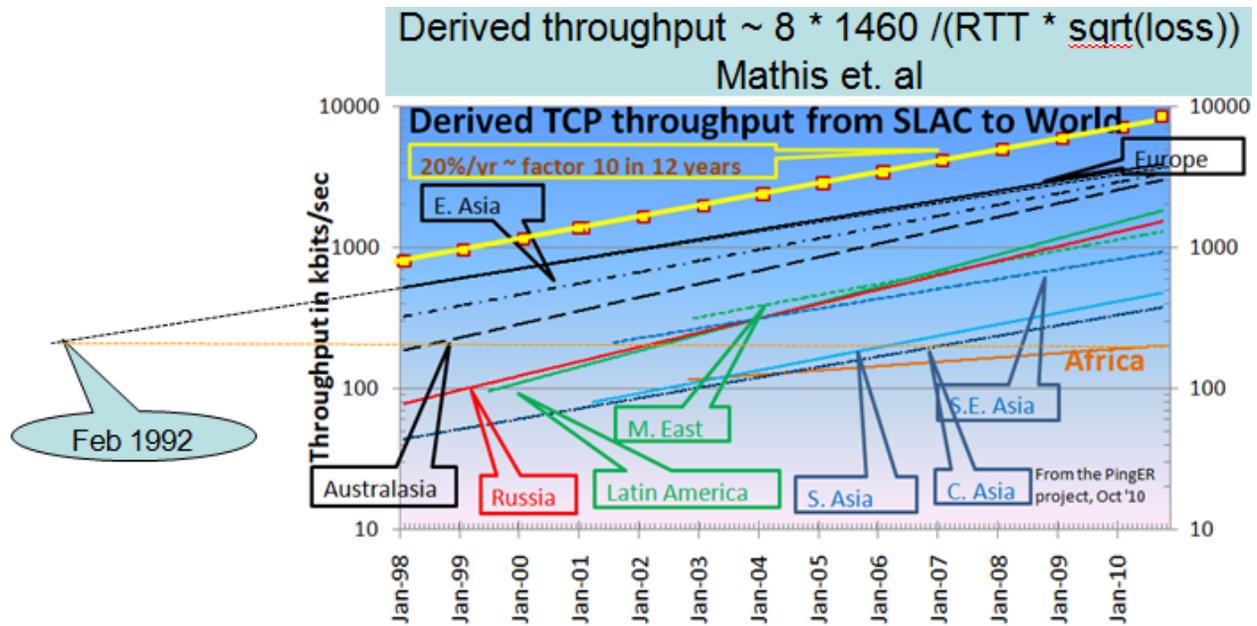


Figure 3: Derived Throughput kilobits per sec from SLAC to the World (since the throughputs in this graph are not normalized we have not shown N. America). The yellow line is to help show the rate of change. If one extrapolates Europe's performance backwards to February 1992, it intercepts Africa's performance today.

The improvements (increased throughput) are seen to be exponentially increasing at roughly 20% /year (the yellow line). Looking at the data points in more detail, one can see the East Asian and Australasian trends catching up with Europe. Russia, Latin America and the Middle East are about 5-6 years behind Europe but are catching up. South East Asia is about 9 years behind Europe and keeping up. South Asia and Central Asia are about 12-14 years behind Europe and also keeping up. Africa is ~ 18 years behind Europe and even worse has been falling further behind. If one extrapolates the trend lines for Africa and Europe to 2020 then at the current rate, Africa's throughput will be 30 times worse than Europe's.

The MOS is a function of the RTT, the jitter and the loss. The defined values of the MOS are: 1= bad; 2=poor; 3=fair; 4=good; 5=excellent. A typical range for Voice over IP is 3.5 to 4.2. Shown below are the MOS values derived from the PingER data.

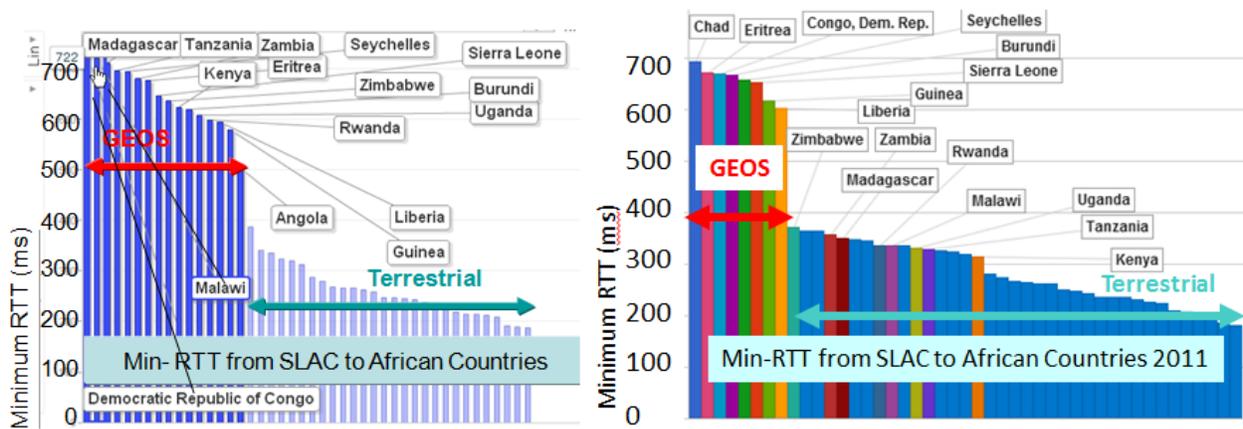


Figure 5: Minimum RTTs measured from SLAC to African countries, 2009 (left), 2011 (right)

It is seen that there is a step change (350ms to 500ms) in minimum RTT as one moves from right to left. The countries to the left are connected to the Internet by geo stationary satellites (GEOS) that rotate at the same rate as the earth and so appear stationary when observed from the earth. GEOS orbit at an altitude of 22,236 miles. The RTT between such a satellite and the earth is about 400ms. The maps below show which countries were connected by GEOS (those in red) in 2008 and at the end of 2010.

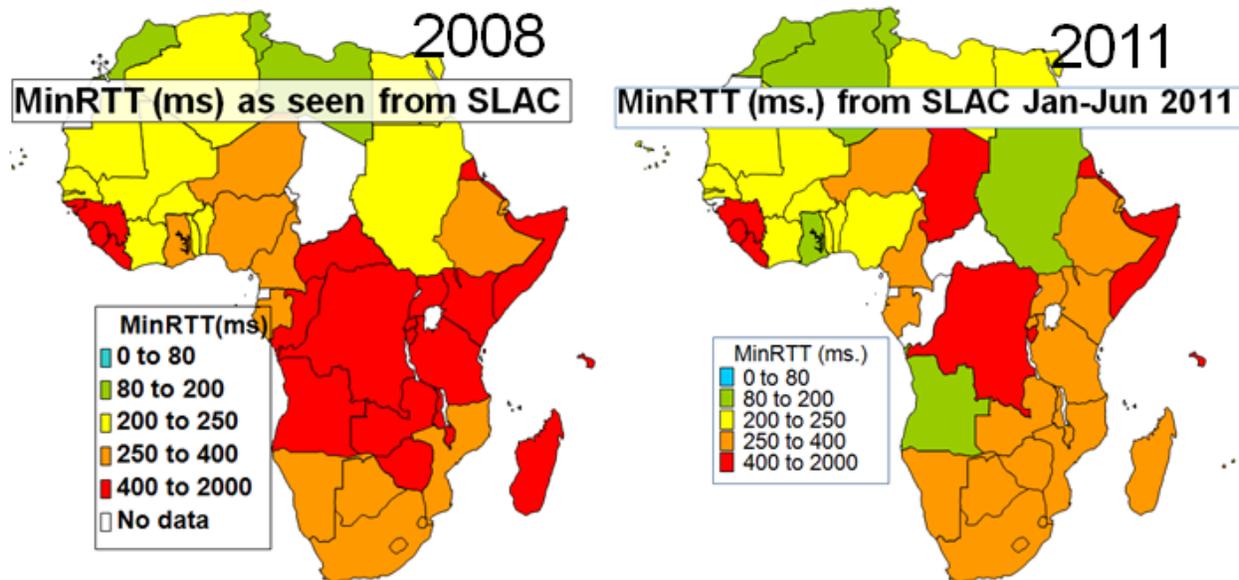


Figure 6: Map of minimum RTTs measured from SLAC to African countries in 2008 and 2011

It is seen that several southern and eastern African countries have moved from using GEOS links to using terrestrial links. This is a result of the installation of five submarine fibre-optic cables being installed along both the east and west coasts of Africa in 2009 and 2010^{viii} driven partially by the communication needs of 2010's Soccer World Cup in South Africa. Prior to this Africa had two terrestrial fibre optic cable links through the Mediterranean and a single fibre optic cable link down the west coast of Africa. An

example of the effect on the RTTs for a host in Kenya seen from SLAC can be seen below.

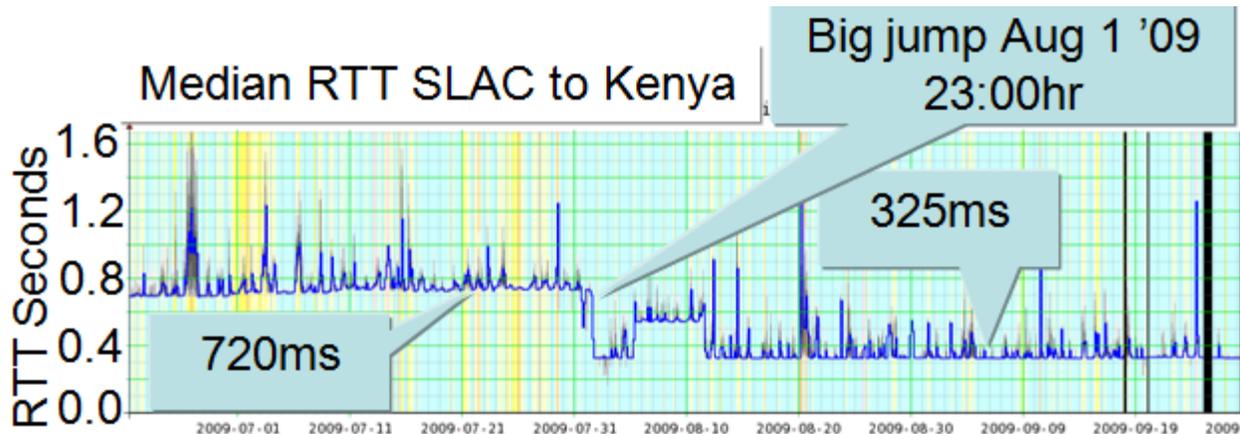


Figure 7: Time series of the median RTT seen from SLAC to a host in Kenya for 120 days in June, July and August 2010. The background color represents the loss, light blue is no loss, light yellow is < 5% loss, and darker yellow is < 10% loss.

The sudden drop in RTT by over a factor of 2 as the international connection was changed from a GEOS to the terrestrial submarine fibre optic cable on August 1st 2009 is clearly evident. At the same time the measured derived throughput improved by a factor of 3. Not only did the new cables increase the international capacity for Africa by 14 times between 2006 and 2010, and bring coverage to the east coast of Africa, but they also provide competition. As a result of this completion, it is speculated that with new competition the megabit-per-second cost would fall from its level in 2009 of over \$1600, to below \$300 by 2011. This would still be far more expensive than Internet connectivity in major developed countries, but it would be a fraction of the cost in 2009^{ix}.

Besides bringing terrestrial connectivity to major coastal cities and the capitals of many African countries with coastal access, the connectivity has spread to landlocked countries such as Uganda and through Uganda to Rwanda using existing fibre optic landlines^x.

The next challenge is the last mile especially for rural communities. The main contenders appears to be: wireless techniques such as microwave, and cell phones; Low Earth Orbiting Satellites such as proposed in 2009 by Google who signed up with Liberty Global and HSBC in a bid to launch 16 LEOS satellites, to bring high-speed internet access to Africa^{xi}; and hydrogen-filled weather balloons, serving in effect as satellite substitutes, at an altitude between 80,000 and 100,000 feet^{xii}.

Internet performance and economic development

The improvement in Internet performance for Africa is critical given its key to competitiveness in today's knowledge economies. Compared to the industrial age one can start a business in one's home with just a simple Internet connection and without the need for roads, railways, canals or access to raw materials such as iron, coal or oil.

A 2010 World Bank report said that for every 10 percentage-point increase in high-speed Internet connections there is an increase in economic growth of 1.3 percentage points^{xiii}. It is thus of interest to

look at the relationship between the Internet performance and development indexes such as the United Nations' Human Development Index or the more technically oriented Information and Communications Technology (ICT) Development Index (IDI)^{xiv} from the ITU. The IDI provides a measure the progress a country is making to become an information society. The IDI is based on ICT readiness + usage + skills, where

- readiness is based on access to infrastructure such as phone (cell & fixed) subscriptions, international bandwidth, %households with computers, and % households with Internet access;
- usage is a function of the % of the population that are Internet users, the % that are mobile users and the % with fixed broadband access;
- the skills are measured by literacy and the % of the population with secondary and tertiary education.

The IDI is in the range 1-10. Its publication was in 2010, and was based on data from 2008. The figure below shows a 'bubble chart' of the PingER derived throughput for 2010 versus the IDI. Each bubble represents a country. The size of the bubble is proportional to the country's population. The color of the bubble represents the region the country is in.

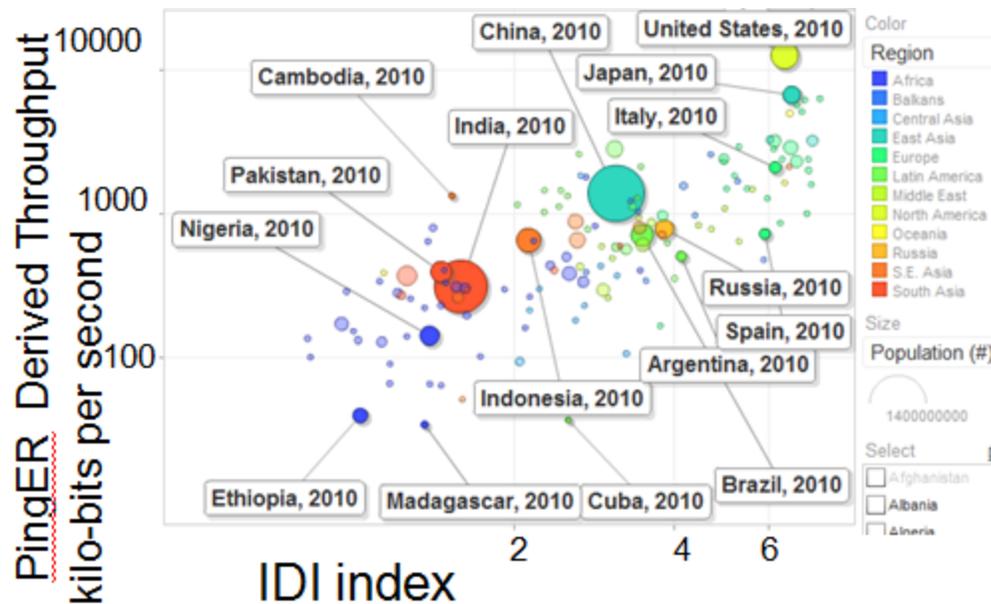


Figure 8: PingER derived throughput versus the IDI

It is seen that there is a positive correlation between the derived throughput and the IDI. This is especially true for the more populous countries (bigger bubbles). Roughly speaking a factor of 10 increase in Internet performance is equivalent to a factor of 3 improvement in the IDI. The PingER data is a useful complement to the IDI especially since the measurement of PingER's data is automatic, does not require an army of data gatherers and statistical analysts, it does not rely on sources that potentially may be politically motivated, and the information is updated on a monthly basis.

National Research and Education Networks (NRENs) in Africa

In addition to the increase in international bandwidth to Africa brought by the new submarine cables, many African countries have or are successfully creating National Research and Education Networks (NRENs) to gather together the universities into an organized bargaining unit and use common standards for their Internet access. Figure 9 below shows the state of NRENs in Africa in 2012.

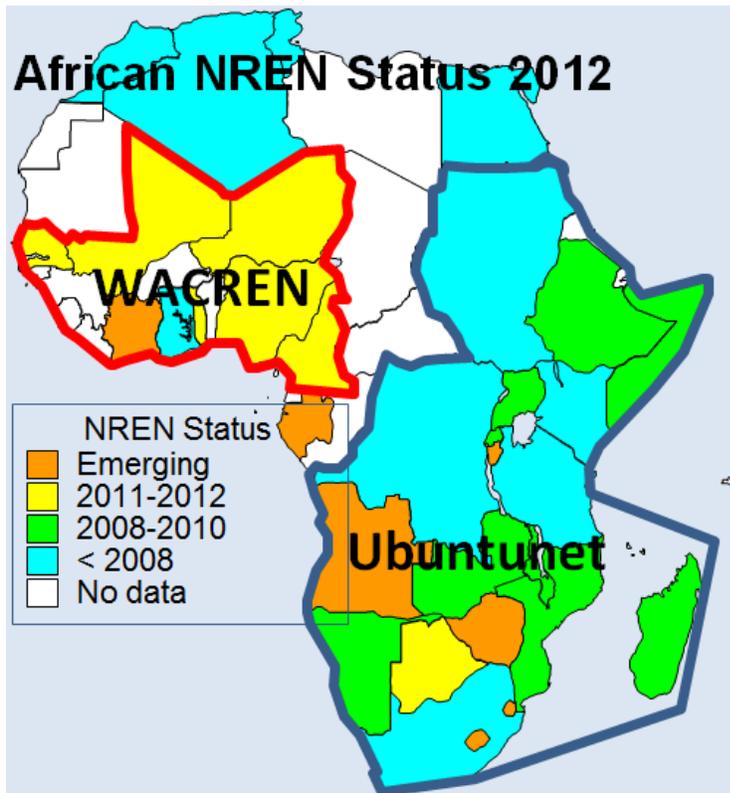


Figure 9: NRENs in Africa in October 2010.

The goal of the UbuntuNet Alliance^{xv} in Eastern and Southern Africa is to “capitalize on the emergence of optical fibre and other terrestrial infrastructure opportunities and thus become the Research and Education Network (REN) backbone of Eastern and Southern Africa. Tertiary education and research institutions throughout the rest of the world are connected to the Internet using fast low-cost fibre. This gives them a huge research and learning bonus”. In January 2008 UbuntuNet linked at 1 Giga-bits per second to the pan-European network GÉANT (the pan European research and Education Network) to provide connectivity for two Sub-Saharan African NRENs, TENET (South Africa) and KENET (Kenya)^{xvi}. Now, in September 2012 June, there are two 10 Giga-bits per second links between UbuntuNet and GÉANT^{xvii}. The North African NRENs are part of the Arab States Research and Education Network (ASREN)^{xviii} connected to GEANT via EUMEDCONNECT^{xix}. The West and Central Research and Educational Network (WACREN)^{xx} is working with UbuntuNet and GEANT as part of the AfricaConnect Initiative. There are also plans^{xxi} to connect ASREN, WACREN and UbuntuNet via interconnections in Khartoum or Nairobi, Mauritania, and the Democratic Republic on the Congo.

Conclusion

Despite the dire state in the past of Internet performance in most of Africa, there are reasons for hope for improvement. Some of these are mentioned above, others include:

- the attraction to businesses of a enormous, youthful, untapped market. Examples of this interest, just in the telecommunications field, include:
 - the enormous investments by telecommunications companies to install six new Sub-Saharan submarine cables between 2009 and 2012;
 - the investment of over \$14.75M by the EU to connect up GEANT to UbuntuNet;
 - the technical trials for 4G cell phone technologies involving a \$141.2 million three year strategic partnership between Safaricom, Kenya’s largest cellular company, and the Chinese firm Huawei Technologies^{xxii}
- the ability to leapfrog technologies such as:
 - personal computers in favor of tablets such as iPADS, smartphones and netbooks;
 - wired connections in favor of wireless

As Eric Hersman, co-founder of Ushahidi put it: *The default device for Africa is a basic mobile, and “SMS technology is now.” ... “the phone number trumps the email address”, Technology is developing at such a rate that, “If you blink, you’ll miss it.” Although basic mobile phones are still the most ubiquitous and thus useful devices in Kenya, there is a predicted 843% mobile internet growth by September 2011 compared to 1 year previous. 60,000 Android-powered IDEOS phones have now been sold in Kenya.*

ⁱ <http://bigthink.com/ideas/21084>

ⁱⁱ Over 50% of Internet Service Providers (ISPs) in Africa are government controlled

ⁱⁱⁱ TeleGeography in a 5 October 2010 press release titled “Global Internet Traffic Growth Remains Strong in 2010”,

^{iv} <http://www.arp.harvard.edu/AfricaHigherEducation/Factoids.html>

^v <http://www.internetworldstats.com/stats.htm>

^{vivi} http://www.worldmapper.org/animations/internet_users_animation.html

^{vii} <http://www.iepm.slac.stanford.edu/pinger/site.html>

^{viii} <http://manypossibilities.net/african-undersea-cables/>

^{ix} <http://www.helium.com/items/1941257-growth-of-the-internet-in-africa?page=2>

^x <http://www.ubuntunet.net/fibre-map>

^{xi} <http://gigaom.com/2008/09/09/google-invests-in-satellite-based-internet-startup/>

^{xii} http://www.internetevolution.com/author.asp?section_id=694&doc_id=178131&

^{xiii} <http://www.infodev.org/en/Article.522.html>

^{xiv} <http://www.itu.int/ITU-D/ict/publications/idi/2010/index.html>

^{xv} <http://www.ubuntunet.net/>

^{xvi} <http://www.erina4africa.eu/news/traffic-increase-drives-upgrade-to-ubuntunet-and-geant-interconnection>

^{xvii} <http://www.geant.net/pages/home.aspx>

^{xviii} <http://www.asrenorg.net/about.html>

^{xix} <http://www.eumedconnect.net/>

^{xx} <http://www.wacren.net/>

^{xxi} <http://www.ubuntunet.net/may2012#gambia>

^{xxii} <http://www.balancingact-africa.com/news/en/issue-no-520/internet/safaricom-and-huawei/en>