Research Notes

Scientific Measure of Africa's Connectivity

M. Zennaro
mzennaro@ictp.it
Abdus Salam International
Centre for Theoretical Physics
Trieste, Italy

E. Canessa
canessa@ictp.it
Abdus Salam International
Centre for Theoretical Physics
Trieste, Italy

K. R. Sreenivasan
director@ictp.it
Abdus Salam International
Centre for Theoretical Physics
Trieste, Italy

A. A. Rehmatullah
aziz@ictp.it
Stanford Linear Accelerator
Stanford, California, U.S.A.

R. L. Cottrell
cottrell@ictp.it
Stanford Linear Accelerator
Stanford, California, U.S.A.

Abstract

Data on Internet performance and the analysis of its trend can be useful for decision makers and scientists alike. Such performance measurements are possible using the PingER methodology. We use the data thus obtained to quantify the difference in performance between developed and developing countries, sometimes referred to as the "digital divide." Motivated by the recent interest of G8 countries in African development, we particularly focus on the African countries.

Introduction

Science and Technology (S&T) are critical for social and economic development within a nation and between developed and developing countries (Special Report 2005).[1] Fortunately, there are signs that S&T are increasingly becoming a part of the agenda of the international community and policy makers within developing countries (Weerawarana and Weeratunge 2004).

Achieving scientific development will depend in part on increased cooperation between scientists from the North and the South, including setting up networks of researchers and institutions. Modern collaboration requires the use of the Internet. Network connection can be used for large-scale scientific data transfer, real-time collaboration, or access to scientific literature. Unfortunately, the network connection in developing countries is marginal at best (Cerdeira et al. 2003). There is much discussion on the digital divide but very little concrete, current data quantifying it. To set expectations for the quality of connectivity, it is critical to monitor the performance of the Internet—for instance, in remote areas of Africa. Monitoring is important to set realistic expectations, understand where upgrades are needed, and provide troubleshooting information. Data on the trends of Internet connectivity can be useful even for decision makers within the G8 when they discuss further issues related to resource allocation on debt relief for Africa.

The PingER project

Performance measurements are possible through the PingER (Ping End-to-end Reporting) project (Matthews and Cottrell 2000), developed by the Internet End-to-end Performance Measurement (IEPM) group at the Stanford Linear Accelerator Center (SLAC). The project started monitoring end-to-end network performance for high-energy physics (HEP) experiments in the mid-1990s. The technique involves sending out a pulse that is echoed

1. Address all correspondence to sdu@ictp.it.

Information Technologies and International Development
Volume 3, Number 1, Fall 2006, 55–64

back at the remote host and timing the round-trip time (RTT). The monitoring has low network impact and can be used for hosts with especially poor connectivity. The resulting monitoring information is in the public domain. Today more than 750 remote sites are being monitored worldwide, more than 45 of them in Africa (PingER data for Africa 2006).

PingER is the name given to the IEPM project to monitor the performance of Internet links. The main mechanism used is the Internet Control Message Protocol (ICMP) echo mechanism, also known as the Ping facility (Tanenbaum 2003). ICMP packets are special Internet protocol control messages that are used to send network information between two hosts. The Ping facility allows sending an echo request packet of a user-selected length to a remote node and having it echoed back. The response provides useful information such as the RTT it took to receive a reply, the variability of response time (jitter), the number of packets that were lost, whether the packets were received in order, whether the remote site is reachable (no response for a succession of pings), and so forth. Ping usually comes preinstalled on almost all platforms. The server (i.e., the echo responder) runs at a high priority (e.g., in the kernel on Unix) and so is more likely to provide a good measure of network performance than a user application.

PingER has a very low network impact of less than 100 bits/second for each monitoring-remote host and can be set to less than 10 bits/second for hosts with especially poor connectivity. It also provides end-to-end network information (as opposed to how well some part of the network such as a backbone is working) and so should be closely related to end-user perceptions.

There are more than 750 remote sites in more than 123 countries being monitored by in excess of 30 active monitoring Sites in 14 countries. These countries contain more than 80% of the world's population and more than 99% of the online users of the Internet. Most of the hosts monitored are at educational or research sites. PingER has historical data going back to January 1995, so a wealth of trend information is available (Cerdeira et al. 2003).

Packet Loss is a good measure of the quality of the link for many applications. Loss is typically caused by congestion, which in turn causes queues (e.g., in routers) to fill and packets to be dropped. Losses may also be caused by the network delivering an imperfect copy of the packet. This is usually caused by bit errors in the links or in network devices. When we get a zero-packet-loss sample (a sample refers to a set of *n* pings[2]), we refer to the network as being quiescent or nonbusy. We can then measure the percentage frequency of how often the network was found to be quiescent. A high percentage is an indication of a good network. This is particularly true for interactive applications, such as video conferencing and audio chat, which require a low packet-loss percentage and is also true for bulk data transport applications, because losses cause long delays for timeouts and so forth. The loss levels that we use to describe the link quality are the following: 0–1% of packet loss is good, 1%–2.5% is acceptable, 2.5%–5% is poor, 5%–12% is very poor, and greater than 12% is regarded as bad. Our observation is that above 4–6% packet loss video conferencing becomes irritating and speakers of the nonnative languages are unable to communicate. Packet loss greater than 10–12% is an unacceptable level of back-to-back loss of packets and long timeouts, TCP connections start to break, and video conferencing is unusable. These values are reinforced by ITU studies of the impact of loss on voice conversations (ETSI 2006).

The RTT is another indicator of the quality of a link. Unlike packet loss, however, where it is possible to reduce losses to zero, it is impossible to reduce the RTT to less than the time taken for information to travel the distance along a fiber or copper cable or wirelessly (e.g., to a satellite and back). In addition to the cable or wireless delays, the measured RTT is the time taken for the packet to be accepted by the router interface, the delay caused by the queuing, and the time taken for the packet to be passed to the outbound interface. The minimum RTT thus usually indicates the length of the route taken by the packets, the number of hops, the line, and router speeds. Increasingly, as speeds of routers and links increase, the main effect on minimum RTT is the aggregate distance of the links in the path. Changes in the minimum RTT can be an indication of a route change. The major effect of poor response time is felt on interactive applications such as telnet or "packetized" video or voice, where even

2. The default value for n is 10.

a moderate delay (e.g., ITU studies show that RTTs of more than 250–300 milliseconds result in interactive voice participants having difficulty to know when the other person has finished speaking) can cause severe disruption. Applications that do not require such a level of interactivity, such as e-mail and web browsing, may appear to perform well even with high delay.

Uses of the PingER Data and Decision Makers

During the past 6 years, the information gathered by the PingER project has been used in several ways. For example, it has been used to set expectations, track network infrastructure changes, and illustrate the need for upgrades to a network. Not all sites that are located in a developing region see the negative effects of the digital divide. If one site can attain credible connectivity, other sites in that region should be able to have better connectivity as well. PingER has also been used to illustrate the digital divide. A PingER brochure was also prepared for the WSIS (World Summit on the Information Society) Second Phase Tunis 2005 conference to throw light on the connectivity problems prevailing in Africa. Other practical uses of PingER include selecting Internet service providers (ISPs) and monitoring the accessibility to network changes and upgrades, and their effects on connectivity.

On the basis of the presentation of the PingER findings, a successful recommendation can be made to people in charge of policy and funding to increase the bandwidth. For example, most recently, a major upgrade for the last mile connectivity for the Pakistan National University of Sciences and Technology (NUST) Institute of Information Technology (NIIT) is being installed based on results from PingER. PingER results derived using the Mathis formula[3] to provide an indicator of end-to-end quality (based on the RTT and loss) correlate positively with economic and development indicators developed by the UNDP (United Nations Development Programme) and the ITU (International Telecommunications Union). These indicators are intended to help policy makers define technology strategies. A full analysis of such indicators and how they correlate with PingER results has been reported by ICFA SCIC (2006). In particular we note here (see figure 1) the correlations between PingER measurements and results for the UNDP technology achievement index (TAI 2001), which is a major international assessment to capture how well a country is creating and diffusing technology and building a human skills base (i.e., a community's potential to participate in the networked world of the future). The TAI focuses on four dimensions of technological capacity that are important for reaping the benefits of the network age: creation of technology (e.g., number of patents granted and receipts of royalties and license fees); diffusion of recent innovations (e.g., diffusion of the Internet and export of high- and medium-technology products as a share of all exports); diffusion of old innovations (e.g., telephones and electricity), and human skills (e.g., mean years of schooling and gross enrollment ratio of tertiary students enrolled in science, mathematics, and engineering).

In figure 1 PingER throughputs measured from North America are plotted against the TAI. Some of the outlying countries are identified by name. Countries at the bottom right of the graph may be concentrating on policies of Internet access for all, whereas countries at the upper right may be focusing on excellent academic and research networks to provide leadership. ICFA SCIC (2006), also features dedicated sections on Africa with the aim of quantifying the digital divide with a vision to educate the HEP community and other scientists on the state and progress in the continent. The U.S. Internet2 group has set up a special interest group on "hard to network places" that includes not only academic and research people but also people from the World Bank and several funding agencies.

As a troubleshooting tool, PingER has been used to discern whether a reported problem is related to networks, identify the time at which the problem has started, decide whether it is still occurring, and provide quantitative analysis for ISPs.

3. The derived TCP throughput is calculated using the formula given by Mathis et al. (1997) relating throughput to packet loss and RTT. The formula is TCP throughput ~ MSS/(RTT × sqrt[loss]), where MSS is the maximum segment size which we take as 1460 bytes. This formula assumes the losses and RTT as measured by TCP. These may be different from those measured by Ping since, for example, TCP provokes loss as part of its congestion avoidance mechanism. Thus, our use of this formula is to provide a convenient measure of performance quality by combining the ping RTT and loss measurements.

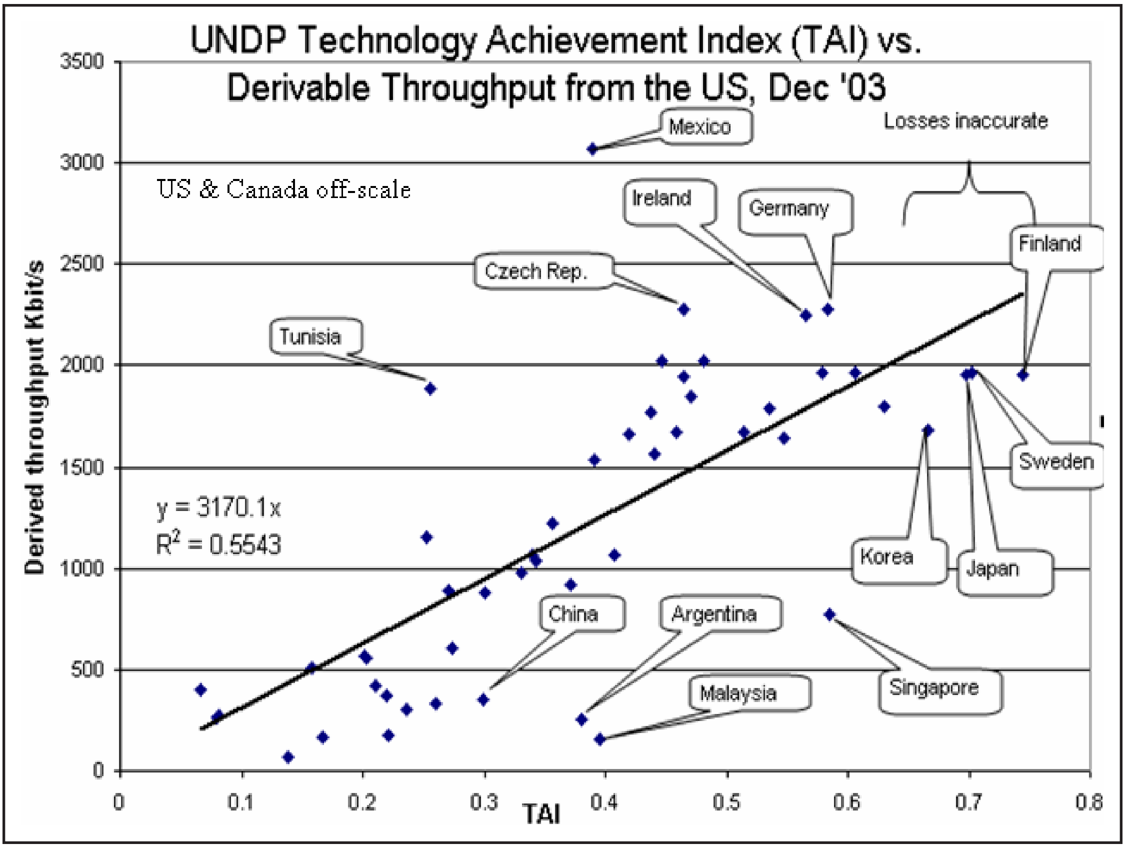

Figure 1. PingER throughputs measured from North America versus the UNDP TAI.

The PingER data can also be used to select universities from developing countries for remote collaboration programs. By using PingER to measure the packet loss and RTT, it is possible to provide expectations (Mathis et al. 1997) on the performance of bulk data transfers and other applications. In the case of real-time collaboration, by comparing the results from PingER with various recommendations for loss, RTT, and jitter, together with the perceptions of voice quality from the users, one can determine how well VoIP and other interactive applications might work between various pairs of sites (ITU-T 1999).

Results of PingER monitoring in Africa

Coverage

As can be seen in figure 2, the PingER project monitors more than 40 African institutions in countries that between them contain more than 80% of the population of Africa. To obtain hosts to monitor at remote sites in Africa, contacts were found by sending e-mails to colleagues especially in the International Committee for Future Accelerators (ICFA) and the ICTP's eJDS (Electronic Journals Delivery Service). Sometimes these e-mails resulted in further referrals or required extended explanations. It was then checked to determine whether the designated host was accessible to Pings, was truly located in Africa[4] and then entered in the PingER database. Typically about 75% of the contacts eventually resulted in a remote host to monitor successfully. Sometimes it took many months to conclude agreements.

All the African sites are monitored from SLAC in the USA. Although the minimum RTT depends on the distance between the monitoring host and the remote host, for a given remote host there is little difference in the losses measured from different

4. Several of the initially chosen hosts were web servers that had proxy servers located outside Africa.

Figure 2. Monitored African sites and the monitoring site in South Africa (red dot). Most of the hosts monitored are at educational or research sites.

monitoring hosts. This indicates that the common bottleneck in most cases is closely associated with the remote site.

The increase in the number of African sites monitored by PingER in the past six years can be seen in figure 3 together with the losses seen by the site. This reflects the dramatic increase in interest in Africa following various meetings focusing on Africa, especially those organized by the ICTP (2002), the ICFA Standing Committee on Regional Interconnectivity (SCIC 2006), and the WSIS.

Quality

The new millennium is beginning to see significant advances in Africa's quest for greater connectivity (International Workshop on African Research and Education Networking 2005). Nevertheless, although a substantial increase in the rate of expansion and modernization of networks is taking place, the ITU statistics (ITU Digital Access Index 2006) on teledensity show that although there are 57.3 Internet users per 100 inhabitants in Sweden, 57 in the United States, and 34.7 in Italy, there are just 0.5 in Mali and 0.2 in Niger. The Internet tariff for the same type of connection is 1.1% of the Gross National Income in Sweden and in the United States and 1% in Italy, whereas it is 289% in Mali and 683% in Niger. The same differences are reflected by the Internet performance. Figures 4a and 4b show the Internet performance quality using the derived throughput from SLAC in the United States and from the Conseil Européen pour la Recherche Nucléaire (CERN) in Europe to world regions. We show both graphs to help distinguish the effects of distance, because the derived throughput is inversely proportional to the RTT and thus the distance between the regions. For example, the fact that in figure 4b the performance from Europe to Europe is much better than for Europe to North America is due to the RTT from CERN to North America being much greater than CERN to European sites.

It is clear that the Internet performance to Africa is slower by a factor of 50, on average, when compared to that between Europe and North America. What can also be seen from these graphs is that

- In the long term, performance to all regions is improving;
- For the developed regions performance is improving by roughly a factor of 10 in 5 years;
- Performances in developing regions are a factor of 5–20 times worse than that in developed regions;
- Performance to developing regions is typically on a par with what was seen 2–7 years ago in developed regions; and

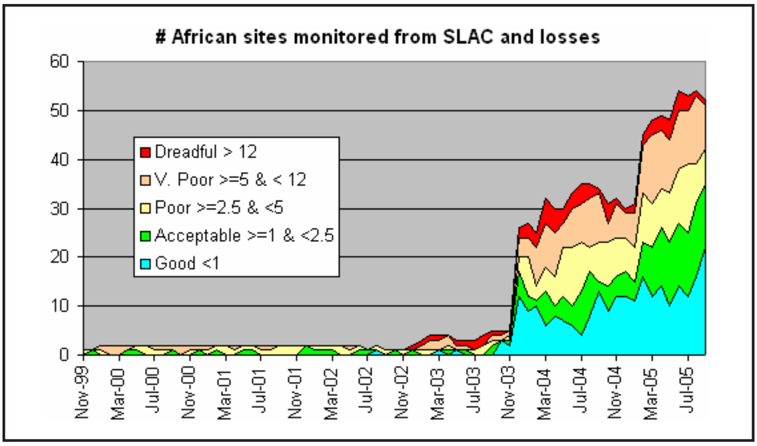

Figure 3. Number of sites in Africa monitored by PingER from SLAC, also showing the packet loss percentages.

- Africa is not catching up with the developed regions, unlike Russia and Latin America.

When Africa's performance is analyzed by region as in figure 5, North Africa performs well, whereas East Africa is falling behind. In North Africa, Algeria and Morocco perform very well, benefiting the entire region. From the existing results, it is apparent that most African regions have poor to bad connectivity. In fact many sites supporting hundreds of users appear to have smaller throughput than many homes with DSL, cable, or dial-up modems in developed countries. Even within the same region, there are differences of more than an order of magnitude in performance between different countries.

Routing

In August 2005, the PingER project established a monitoring station at the TENET site in Ronderbosch,[5] South Africa, so now there are measurements from within and between African countries. The data, at the start of 2006, indicated that for the sites in the 28 African countries monitored by PingER from this South Africa monitoring station:

- The only countries with sites that had direct connectivity (i.e., without going through Europe or North America) to South Africa are Botswana, Zimbabwe, and Tunisia. Even then one of the three sites monitored in Tunisia connected through North America.
- The sites in fifteen countries (Algeria, Benin, Eritrea, Ethiopia, Kenya, Madagascar, Malawi, Mali, Mauritania, Morocco, Mozambique, Niger, Nigeria, Senegal, and Sudan) connected through Europe.
- The sites in five countries (Angola, Cameroon, Egypt, Namibia, and Uganda) connected through North America.
- The sites in two countries (Ghana and Tanzania) connected through Europe or North America.
- The sites in two countries (Burkina Faso and Rwanda) connected through Europe and North America.

Not only does this situation drastically increase the RTT, but also utilizes expensive (for African countries) intercontinental links since the Africa Internet Service Provider typically has to pay 100% of the international carrier cost to get from Africa to the International Backbone Provider (AFRISPA 2006).

Minimum RTT

It is also instructive to assess the fraction of African countries that have sites connected to South Africa by geosynchronous satellites. This is illustrated in figure 6. As might be expected, the countries with terrestrial connections are those served by the SEAMEWE II/III and SAT 3 cables or sharing a border with South Africa. The impact of a geosynchronous satellite is to add roughly 600 milliseconds to the minimum RTT. It should be noted that the East African Submarine cable System (EASSy 2006) should greatly assist in improving the connectivity of East African countries.

Attention should be paid to the usability of the Internet in view of these performances. If we assume that the same network supports a number of computers that can reach the hundreds in many university labs and that bandwidth requirements are growing every day (both with increases in number of users and increasing use of more bandwidth hun-

5. Many thanks to Duncan Martin and Len Lotz of TENET.

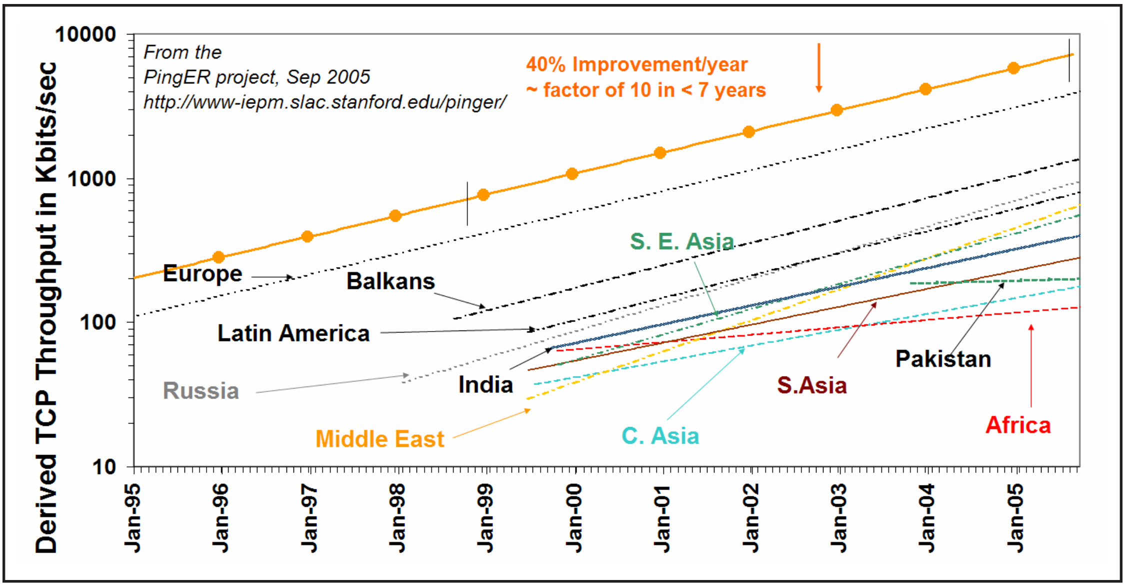

Figure 4a. Throughput performance from SLAC to regions of the world. The lines are fits to exponential functions to help guide the eye.

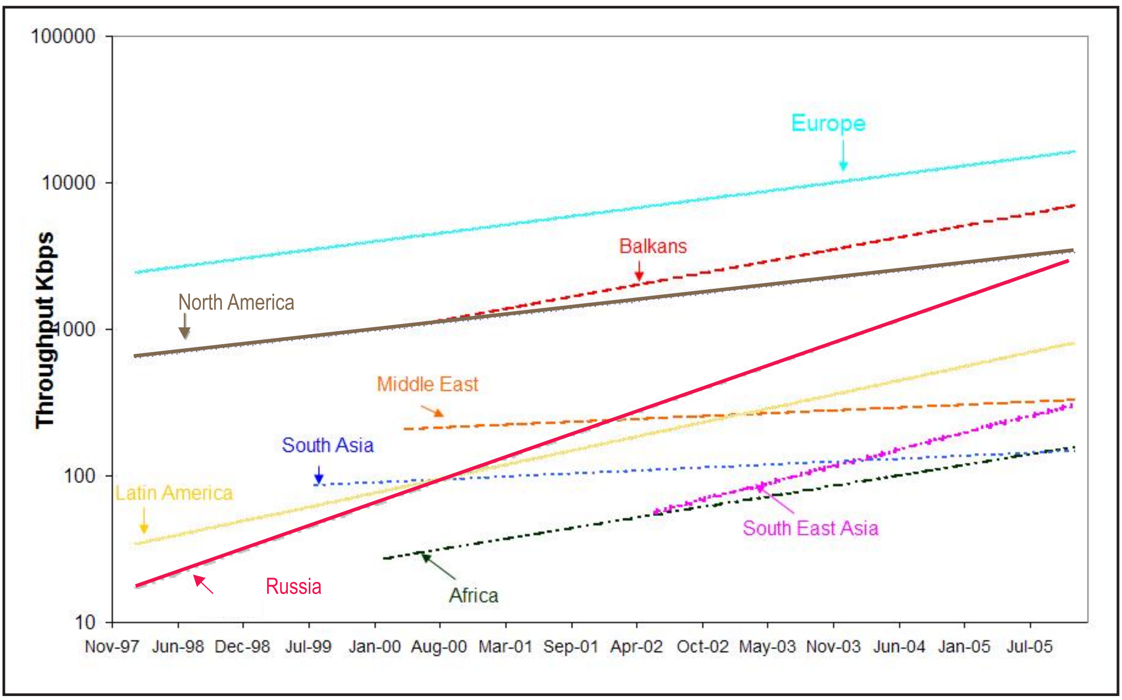

Figure 4b. Throughput performance from CERN to regions of the world.

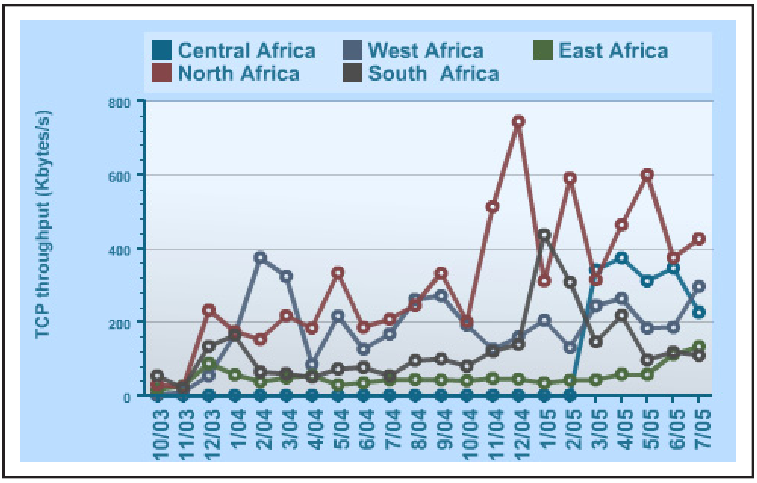

Figure 5. Throughput performance from SLAC to different regions of Africa.

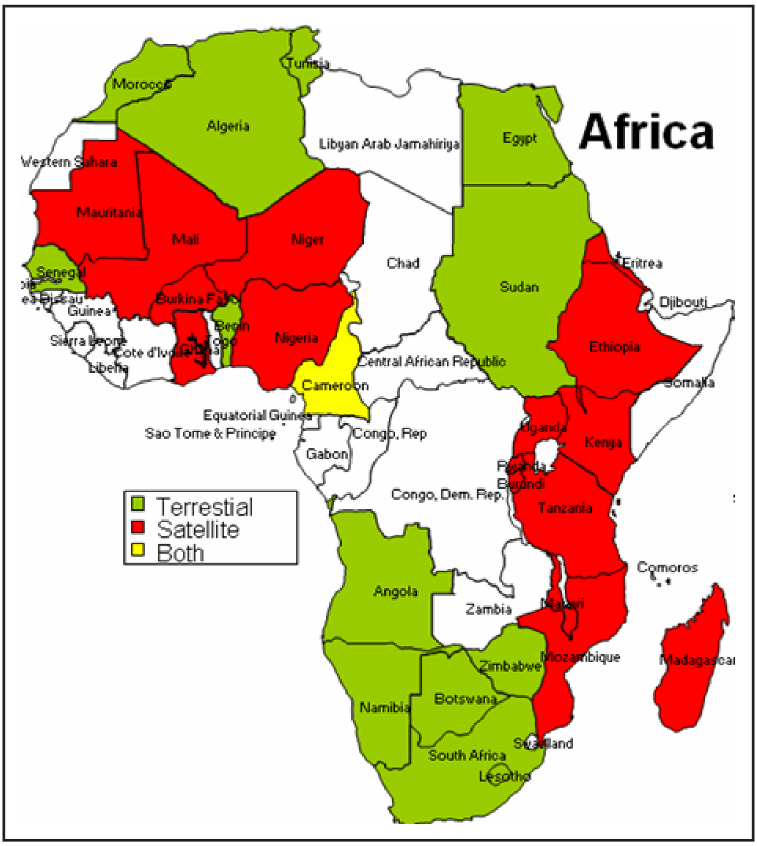

Figure 6. African countries with sites connected to South Africa by geosynchronous satellites and by terrestrial links.

gry applications), it is clear that most Internet applications are increasingly impossible to use. ICTP has developed a "Digital Divide Simulator" (ICTP 2006) that enables scientists with high bandwidth connections to experience what their colleagues from the South experience daily when using the Internet. Other initiatives, such as the eJDS, are being supported by the center to deliver scientific literature to scientists with low-bandwidth connections.

Fragility

PingER sends ten 100-Byte packets/probes for each 30-minute period. When no response is received after ten probes, the node is categorized as down or "unreachable" during that interval. The unreachability is then defined as the number of periods the node is unreachable divided by the total number of periods. For most regions the unreachability is improving by a factor of 5 in 2 years. As seen from SLAC, Europe has the most stable infrastructure, with Russia fast catching up. Sites in South Asia and Africa have roughly a factor of 10 worse unreachability than developed regions such as Europe. Despite being connected via satellite, the sites in Central Asia have better unreachability compared to Africa. The unreachability factor can be attributed to the frequent network outages resulting from weak infrastructure such as power, or lack of adequate backup routes. Another factor is the frequent change in configuration of the hosts or their connections, which make them unresponsive to Ping measurements.

Conclusions

Continual efforts are being made to maintain, upgrade, and expand the PingER deployment. With respect to the remotes sites, the primary effort is to keep them accessible to the monitoring sites. For monitoring sites, it is important to keep the list of remote sites up to date, the monitoring code running reliably, and the data accessible for collection. The data and the reports are made available publicly available via the web (SLAC 2006). Interested users may use the data selection, analysis and display tools (SLAC 2006) to make their own analyses.

Much work needs to be done to extend the monitoring to more of Africa. Even for the countries already monitored, more sites are needed to help avoid anomalous results associated with a single site being used to represent an entire country. To understand where South-to-South scientific collaboration is feasible using the Internet, it is necessary to monitor more African sites from South Africa and to set up monitoring hosts in other African countries besides S. Africa. The authors encourage readers to provide contacts for new countries and sites.

PingER provides a valuable lightweight tool for the active end-to-end performance monitoring of networks around the world. These measurements indicate resource availability to a user. With its continual gathering of Internet monitoring data, it provides extensive quantitative historical and almost real time information on worldwide networks. PingER has shown itself to be useful for providing valuable information quantifying network needs and improvements. Its results can be used by the G8 countries involved in African development, to select the countries that need support to develop their network, and to monitor the effectiveness of the improvements.

References

AFRISPA Initiatives. Retrieved July 2006. http://www.afrispa.org/initiatives .htm.

Aldous, P., et al. 2005. "Special Report: Science & Africa." *Nature* 436: 1146–49.

Canessa, C., H. Cerdeira, E. Fonda, and R. L. Cottrell. 2003. "Developing Countries and the Global Science Web." *CERN Courier* 43: 17–18.

Cottrell, R. L., and A. A. Rehmatullah, eds. 2006. *ICFA SCIC Network Monitoring Report*. http://www.slac.stanford.edu/xorg/icfa/icfa-net-paper-jan06/.

Developing Country Access to On-Line Scientific Publishing: Sustainable Alternatives. Round table meeting, ICTP Trieste, October 2002. http://www.ictp.it/?ejds/seminars2002/program.html.

Digital Divide Simulator. Retrieved July 2006. http://wireless.ictp.it/simulator.

East African Submarine cable system. Retrieved July 2006. http://www .eassy.org.

ICFA SCIC Network Monitoring Report. http://www.slac.stanford.edu/xorg/icfa/icfa-net-paper-jan06.

International Workshop on African Research & Education Networking. 2005. http://event-africa-networking.web.cern.ch/event-africa-networking/.

ITU Digital Access Index. Retrieved July 2006. http://www.itu.int/ITU-D/ict/dai/index.html.

ITU-T SG13 Draft Rec. I.380. 1999. "Internet Protocol Data Communication Service—IP Packet Transfer and Availability Performance Parameters." Geneva, February.

Mathis, M., J. Semke, J. Mahdavi, & T. Ott. 1997. "The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm." *Computer Communication Review* 27, no. 3: 67–82.

Matthews, W., and R. L. Cottrell. 2000. "The PingER Project: Active Internet Performance Monitoring for The *HENP* Community." *IEEE Communications Magazine* 38 (May): 130–36.

PingER Data for Africa. Retrieved July 2006. http://sdu.ictp.it/pinger/africa/.

PingER Tools. Retrieved July 2006. http://www-iepm.slac.stanford.edu/cgi-wrap/pingtable.pl.

Report 1st ETSI VoIP Speech Quality Test Event. Retrieved July 2006. http://www.etsi.org/plugtests/history/doc/22td031.doc.

SLAC PingER Data. Retrieved July 2006. http://www-iepm.slac.stanford.edu/pinger/.

Tanenbaum A. 2003. In *Computer Networks,* pp. 449–50. Englewood, NJ: Prentice Hall.

Technology Achievement Index (TAI) from UNDP. hdr.undp.org/reports/global/2001/en/pdf/techindex.pdf.

Weerawarana, S, and J. Weeratunge. In *Open Source in Developing Countries.* SIDA, 2004.