

Quantifying and Mapping the Digital Divide from an Internet Point of View

R. Les. Cottrell+, Shahryar Khan*+, Jerrod Williams+, Akbar Mehdi*+, Umar Kalim*, Arshad Ali*

*NUST Institute of Information Technology (NIIT), Rawalpindi, Pakistan;
+Stanford Linear Accelerator Center (SLAC), California, USA

Abstract

Quantitative knowledge of the magnitude, extent and trends of the Digital Divide are critical to understand and identify the regions most in need of help, to make rational decisions on how to address the problems and to make cases for executives, funding agencies and politicians. We report on a project (PingER) to measure the Digital Divide from the point of view of Internet performance. The PingER project has been measuring Internet performance since 1995 and with the increased emphasis on measuring and tracking the Digital Divide, it now covers over 700 hosts in over 150 countries that between them contain over 99% of the world's Internet connected population. In this paper we will describe the how PingER works, its deployment, the data analysis, and presentation. We also introduce a new PingER visualization tool (ViPER) that provides a more appealing interactive visualization of the PingER data and also works on mobile PDAs. We will also show results from PingER that illustrate the magnitude, extent and trends for the Digital Divide, and also compare PingER results with some human development and technology indices.

Introduction

The Digital Divide “is the gap between those with regular, effective access to digital technologies and those without”¹. It is vitally important to measure the extent of this gap so one can understand, identify the regions most in need of help, make rational decisions on how to address the problem and catch the attention of funding agencies, politicians etc. to help initiate and fund efforts to bridge the gap. Without such efforts the gap will continue to increase for many regions such as S. Asia and Africa, leading to increased poverty², distrust, political instability etc. There have been several initiatives to make such measurements in terms of human development and technology indices³⁴⁵⁶⁷. In this paper we report on a project to measure the Digital Divide from the point of view of

¹ “Digital Divide”, this definition comes from Wikipedia at http://en.wikipedia.org/wiki/Digital_divide

² “Digital Divide: What It is and Why it matters”, available at <http://www.digitaldivide.org/digitaldivide.html>

³ “Technology Achievement Index”, UNDP Report 2001, available at <http://hdr.undp.org/reports/global/2001/en/pdf/techindex.pdf#search=%22Technology%20Achievement%20Index%22>

⁴ “Human Development Indicators”, UNDP 2005, available at <http://hdr.undp.org/>

⁵ “Network Readiness Index”, available www.cid.harvard.edu/cr/pdf/gitrr2002_ch02.pdf, also see http://www.forbes.com/technology/2005/03/09/cx_0309wefrinking.html

⁶ “Digital Access Index, ITU report 2003, available at www.itu.int/newsroom/press_releases/2003/30.html

⁷ “Digital Opportunity Index”, ITU report 2005, available at <http://www.itu.int/osg/spu/statistics/DOI/index.phtml>

Work supported in part by US Department of Energy contract DE-AC02-76SF00515

Contributed to the Asian Applied Computing Conference 2007, 12/13/2007 to 12/15/2007, Kathmandu, Nepal

Internet performance. We also relate the measures from this project to measurements from the UNDP Human Development Indicators project.

The impact of the Internet on the Digital Divide is growing rapidly, and today's network performance targets will move as network technologies continue to advance in the economically favored regions of the world. In the first half of 2004 for example the number of Internet users in China grew from 6.8 to 78 million⁸. It now tops 100 million, while the worldwide number of Internet users recently topped 1 billion⁹. The US Energy Sciences Network's traffic has been increasing by 100% per year for the last 6 years¹⁰. The traffic flowing through the Amsterdam exchange increased fourfold in 2005¹¹. The Large Hadron Collider network between CERN and the US grew from a 9.6 kilobits/sec satellite link in 1985 to multiple 10Gigabits/sec today. These developments have been paralleled by upgrades in the metro, state, national, and continental core network infrastructures, as well as the key transoceanic links used for research and education. The transition to the use of "dense wavelength division multiplexing" (DWDM) to support multiple optical links on a single fiber has made these links increasingly affordable, and this has resulted in a substantially increased number of these links coming into service. At the end nodes the commoditization of Gigabit and 10 Gigabit Ethernet, new buses, and faster cpus are driving performance higher and costs lower.

All of this adds up to an explosion of opportunities. However, the rapid rate of progress, confined mostly to the US, Europe, Japan and Korea, as well as the major transoceanic routes, threatens to open the Digital Divide between the developed and developing regions further. For example the mean bandwidth per networked computer in Africa is less than 4 kilobits/sec¹², average dialup costs are 8 to 10 times that in the U.S. where the per capita income is more than a factor of ten greater, Africa counts one Internet user in every 250 people compared to the world average of one in 35, and a North American and European average of one in every three people¹³, in Africa one year of Internet access costs more than the average annual income¹⁴.

PingER Project

The PingER project¹⁵ is aimed at measuring the end-to-end network performance of Internet paths. Originated in 1995, it was initially focused on High Energy Physics (HEP)

⁸ See the Feb. and Aug. 2004 Reports by the ICFA Standing Committee on Inter-Regional Connectivity (SCIC) at <http://cern.ch/icfa-scic>, or <http://icfa-scic.web.cern.ch/ICFA-SCIC/docs/Documents/20040820-RepICFA-HN.ppt>

⁹ See <http://www.internetworldstats.com/>

¹⁰ Source: W Johnston (LBNL), ESnet manager

¹¹ See www.ams-ix.net/about/stats/index.html

¹² "Africa tertiary Institutions Connectivity Survey: Preliminary Insights", Sept 2004, available at http://www.atics.info/html/about/Initial_response_report1.pdf

¹³ Mike Jensen, "Africa Internet Status", available at <http://www.africaaction.org/docs00/inet0010.htm>

¹⁴ "High cost of Internet: Bane of African economic growth", available <http://www.vanguardngr.com/articles/2002/features/technology/tec609082006.html>

¹⁵ W. Matthews and R. L. Cottrell, "The PingER Project: Active Internet Performance Monitoring for the HEP Community", IEEE Communications Magazine Vol. 38 No. 5 pp 130-136, May 2002.

collaborations, however, with the new millennium it has increased its focus on the Digital Divide.

Method

Every 30 minutes from a monitoring node, we ping¹⁶ a set of remote nodes with 11 pings of 100 bytes. The nodes to be pinged consist of nodes chosen to be of interest to the monitoring site, plus a centrally maintained list of “Beacon” nodes that are monitored by all monitoring nodes. Information on the end nodes and sites is kept in a central configuration database that among other things keeps the host name, the IP address, the location including the country and region and latitude and longitude. The pings are separated by at least one second, and the default system ping timeout is used. The first ping is thrown away (it is presumed to be slow since it is priming the caches etc.¹⁷). The individual ping times for each set of 10 pings are recorded. This is repeated for ten pings of 1000 data bytes. The use of two ping packet sizes enables us to make estimates of ping data rates and also to spot behaviors that differ between small and large packets (e.g. rate limiting). In general the RTT is proportional to l (where l is the packet length) up to the maximum datagram size (typically 1472 bytes including the 8 ICMP echo bytes).

The data is collected on a regular basis (typically daily) from the monitoring nodes by two archive hosts, one at SLAC the other at Fermilab in Chicago, Illinois, that store, analyze, prepare and via the web make available reports on the results. An overview of the architecture is seen in Fig. 1.

Considerable work is expended to ensure the quality of the data. Tests include:

- anomalous minimum RTTs for a host compared to others in the same region or country are investigated to see if the host is really where we believe it to be (typical causes include a host that has a terrestrial path while another has a geostationary satellite, or a web site with a proxy in another country, or very indirect routing between developing countries);
- has the IP address of a host changed;
- has a host stopped responding for several days;
- is an archive host unable to gather the data from a monitoring host;
- database entries are checked for validity e.g. latitude/longitudes are in range, IP addresses are valid, all hosts have a country and region associated with them, plus a latitude and longitude. Typically the latitude and longitude is that of the city in which the host is believed to be located.

¹⁶G. Kessler, S. Shepard, “A Primer on Internet and TCP/IP Tools and Utilities” RFC 2151, 1997, available at <ftp://ftp.rfc-editor.org/in-notes/rfc2151.txt>

¹⁷ Martin Horneffer in <http://www.advanced.org/IPPM/archive.2/0246.html>

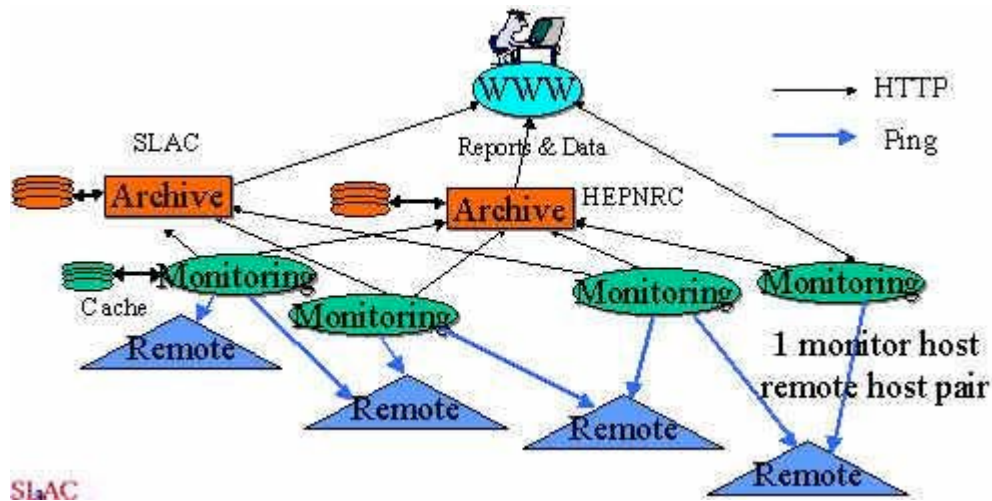


Figure 1: Overall architecture of the PingER project

Analysis

The data is analyzed to extract metrics such as the minimum and average Round Trip Times (RTTs), the losses (a packet is lost if no response is received in the timeout), the variability of the RTTs (jitter), the unreachability (a remote site is unreachable if all 10 pings of a set are lost, and the TCP throughput derived from the Mathis formula¹⁸. This data is aggregated into time bins of an hour, a day, a month and a year, and is available going back to January 1995. It is also aggregated by country and regions, and can be aggregated into affinity groups such as an HEP experiment, sub-regions of the world (e.g. Sub-Saharan Africa), a network (such as ESnet¹⁹ or GLORIAD²⁰) etc.

Deployment

PingER monitors over 700 hosts in over 115 countries from hosts in over 30 countries. This corresponds to over 2200 monitor host / remote host pairs, and about 200,000 ping measurements per day. The deployment of PingER is illustrated in Fig. 2. Considerable effort has been made since 2000 to increase the representation from developing regions. Currently the countries monitored by PingER contain over 99%²¹ of the worlds Internet connected population. Fig. 3 shows the countries monitored by PingER together with which regions of the world PingER assigns them to. The regions are chosen starting from the U.N. definitions²². We modify the region definitions to try and ensure the countries in a region have similar performance.

¹⁸ Mathis, Semke, Mahdavi & Ott, "The macroscopic behavior of the TCP congestion avoidance algorithm" in Computer Communication Review, 27(3), July 1997, provides a short and useful formula for the upper bound on the transfer rate: $Rate < (MSS/RTT) * (1 / \sqrt{p})$, where: *Rate*: is the TCP transfer rate, *MSS*: is the maximum segment size (fixed for each Internet path, typically 1460 bytes), *RTT*: is the round trip time (as measured by TCP), and *p*: is the packet loss rate.

¹⁹ "Energy Sciences Network", available <http://www.es.net/>

²⁰ "Global Ring Network for Advanced Applications Development", available at <http://www.gloriad.org/>

²¹ "Internet Usage Statistics – The Big Picture", available <http://www.internetworldstats.com/stats.htm>

²² "United Nations Population Division World Population Prospects Population database". Available <http://esa.un.org/unpp/definition.html>

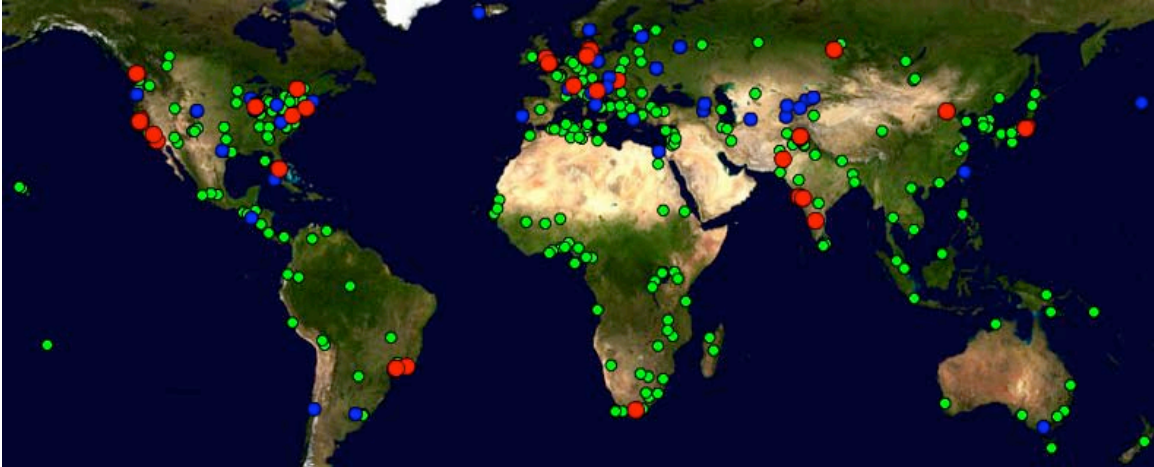


Figure 2: Deployment of PingER monitoring sites (red), beacons sites (blue), and remote sites (green).

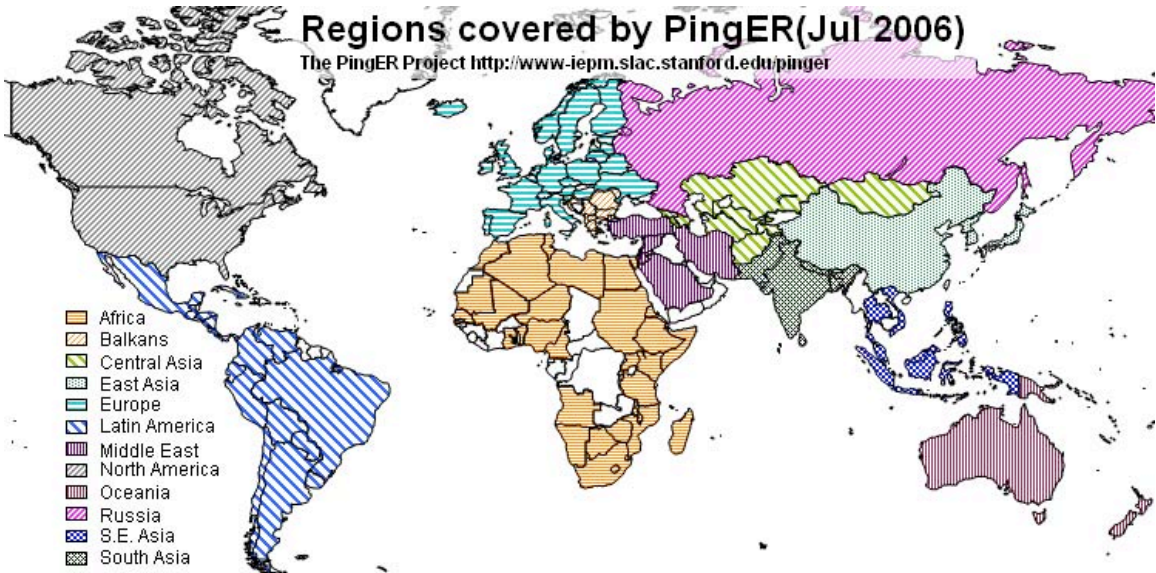


Figure 3: Countries monitored by PingER and the regions they are assigned to. There are no remote hosts monitored in countries in white.

Access/Visualization

The analyzed reports and data and the configuration database information are made publicly available via the web. This enables interested users to download the information and perform their own analyses and generate reports. There are also web forms that enable one to select the metric, the time bin aggregation, the countries/regions and groups for both monitor and remote sites, the ping packet sizes etc. Having made such selections the user can generate tabular and graphical reports. Such user interfaces are valuable for experts, however, they require networking knowledge and are not very visually appealing or interactively oriented. We have therefore developed a more interactive visualization tool known as ViPER²³.

²³ "ViPER (Visualization for PingER)", available <http://www.slac.stanford.edu/comp/net/wan-mon/viper/> and <http://pinger2.niit.edu.pk/~sheryar>

ViPER

ViPER (Visualization for PingER) is a graphically engaging tool that provides an eye-catching view of the PingER's deployment, and the links which are being monitored by PingER. It has a user friendly interface that allows a user to perform interactive analysis on PingER data. It uses a world map to show the geographic location of all the PingER nodes. When the user selects a particular node, detailed information about that node is shown. Also the links are colored according to their performance so the user can compare the performance of various links. ViPER is also valuable to the management of PingER since it quickly identifies PingER database errors i.e. sites with incorrect latitude and longitude.

The user can plot graphs of various metrics such as TCP throughput, minimum RTT and packet loss between a remote monitoring site pair. The graph will show the historical performance of the link. The user can save the analysis information for future reference. A screen shot of ViPER is shown in the Fig. 4.

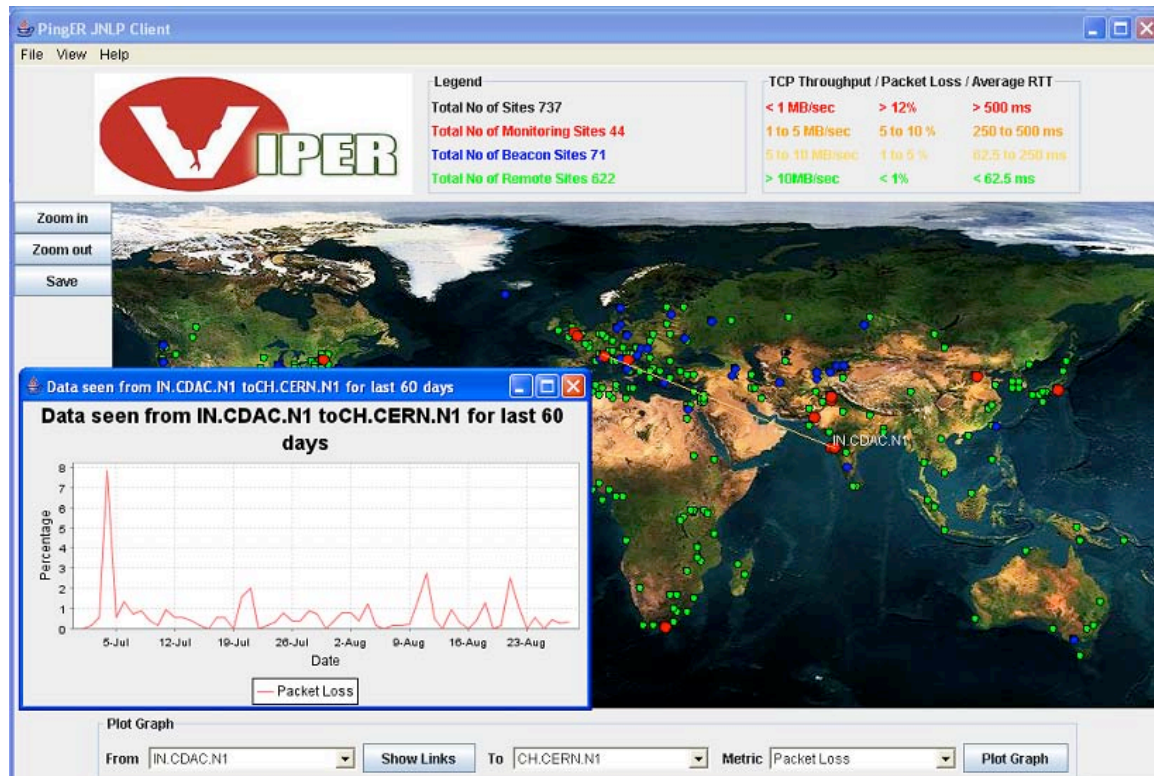


Figure 4: Screen shot of ViPER showing the losses from CDAC India in Mumbai to CERN in Geneva Switzerland

At the bottom of the screen shot is shown the user selection for the monitoring site and remote site, and the metric, also shown on the map in red are the Monitoring sites and in green the remote sites. A colored line from the chosen monitoring site (CDAC in Mumbai) to the remote host (at CERN in Geneva) indicates the quality of the end-to-end path in terms of the chosen metric - loss (the legend of colors is given in the top right).

Moving the mouse over a site will pop up the name of the site or sites if there are multiple sites in the area. Clicking on the site will give detailed configuration information about the site(s). In the case of Fig. 4 the user has also requested to plot the losses for the last 60 days that is shown in the overlaid window.

ViPER is a Java application that has been deployed on JNLP (the Java Network Launching Protocol). Thus the application is launched through a simple mouse click and the users do not need to install the application.

PDA Client for ViPER

Network monitoring information is of great value to Network administrators, Scientists and Researchers. Therefore access to PingER via a PDA may prove to be useful in a variety of circumstances, particularly for brief demonstrations and prompt access.

The PDA Client is a miniature version of ViPER that has been developed for handheld devices, the PocketPC in particular. The PDA client provides ubiquitous access to monitoring information. The code of ViPER has been optimized so that it may run on the limited memory and CPU power of mobile devices. All the features of ViPER are supported by the client application. It also provides zoom-able and drag-able map in order to fit the small screen.

Results

Satellite coverage

Many regions of the world are connected by geostationary satellite links. The minimum RTT between two hosts connected by geostationary satellites is close to 600ms, which is about twice the minimum time needed to propagate a round trip signal between any two places in the world over “as-the-crow-flies” land lines. The minimum RTT therefore provides an effective signature for whether a geostationary satellite is used in the end-to-end path. Fig. 5 shows the minimum RTTs measured from SLAC to the countries of the world in July 2000 and July 2006. In 2006 the main regions using satellite links are East Africa, the Central Asian Republics (currently accessed through the Silk Road project²⁴ with the major ground-station in Hamburg Germany), East Africa (which plans to upgrade to terrestrial lines via the EASSy project²⁵), plus Bangladesh, Nepal, Nigeria, Niger and Ghana. Comparing the turn of the millennium with July 2006, it is apparent that where possible, countries are moving from satellite link to land lines to provide improved performance and latency and reduced costs for countries with access to fibre.

²⁴ “Silk Project Description”, available at <http://www.silkproject.org/project.htm>

²⁵ “East Africa Submarine cable System”, see <http://eassy.org/>

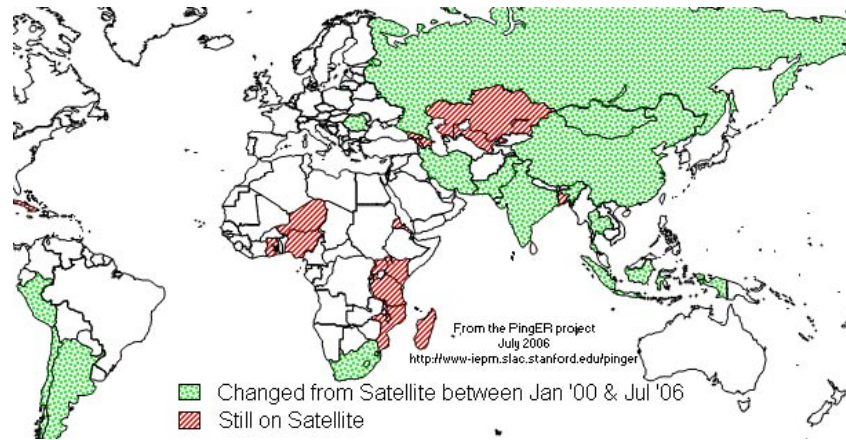


Figure 5: map showing countries with minimum RTTs of over 600ms as measured from SLAC, identifying those that still have this value and those that have improved since Jan '00. Note that a minimum RTT of 600ms typically identifies the end-to-end path as using a satellite link.

Loss Quality

Packet loss is a critical metric for Internet performance since losses typically cause timeouts that can last for several seconds. Losses of 4-6% or more video-conferencing becomes irritating and non-native language speakers become unable to communicate²⁶. Above 10-12% packet loss there is an unacceptable level of back-to-back loss of packets and extremely long timeouts, connections start to get broken, and video conferencing is unusable. The occurrence of long delays of 4 seconds (such as may be caused by timeouts in recovering from packet loss) or more at a frequency of 4-5% or more is also irritating for interactive activities such as telnet and X windows²⁷. A random loss of 2.5% will result in Voice over Internet Protocols (VOIP) becoming slightly annoying every 30 seconds or so. A more realistic burst loss pattern will result in VOIP distortion going from not annoying to slightly annoying when the loss goes from 0 to 1%²⁸. Since according to the Mathis formula TCP throughput for the standard (Reno based) TCP stack goes as $1/\sqrt{\text{loss}}$ it is important to keep losses low for achieving high throughput. To help in relating the losses experienced, we define packet losses of: < 1% as Excellent; between 0.1% and 1% as Good; between 1% and 2.5% as Acceptable; between 2.5% and 5% as Poor; between 5% and 12% as Very poor; and greater than 12% as bad.

Fig. 6 shows the losses measured from SLAC to the world going back to January 1998. It is seen that the number of hosts monitored from SLAC has increased seven-fold with the increasing emphasis on the Digital Divide. It is also seen that despite adding hosts in the less developed regions of the world, the fraction of hosts with good packet loss (< 1%) has increased by about 50%.

²⁶ M. Zennaro, E. Canessa, K. R. Sreenivasan, A. Rehmatullah, R. L. Cottrell, "Scientific Measure of Africa's Connectivity", presented at the International Workshop on African Research and Education Networking, CERN, Switzerland Sept 2005.

²⁷ R. L. Cottrell, W. Matthews, C. Logg, "Tutorial on Internet Monitoring & pingER at SLAC", available at <http://www.slac.stanford.edu/comp/net/wan-mon/tutorial.html>

²⁸ C. Dvorak, "Speech Quality Impact of Random vs. Bursty packet Losses", from ITU-T SG12, 2001.

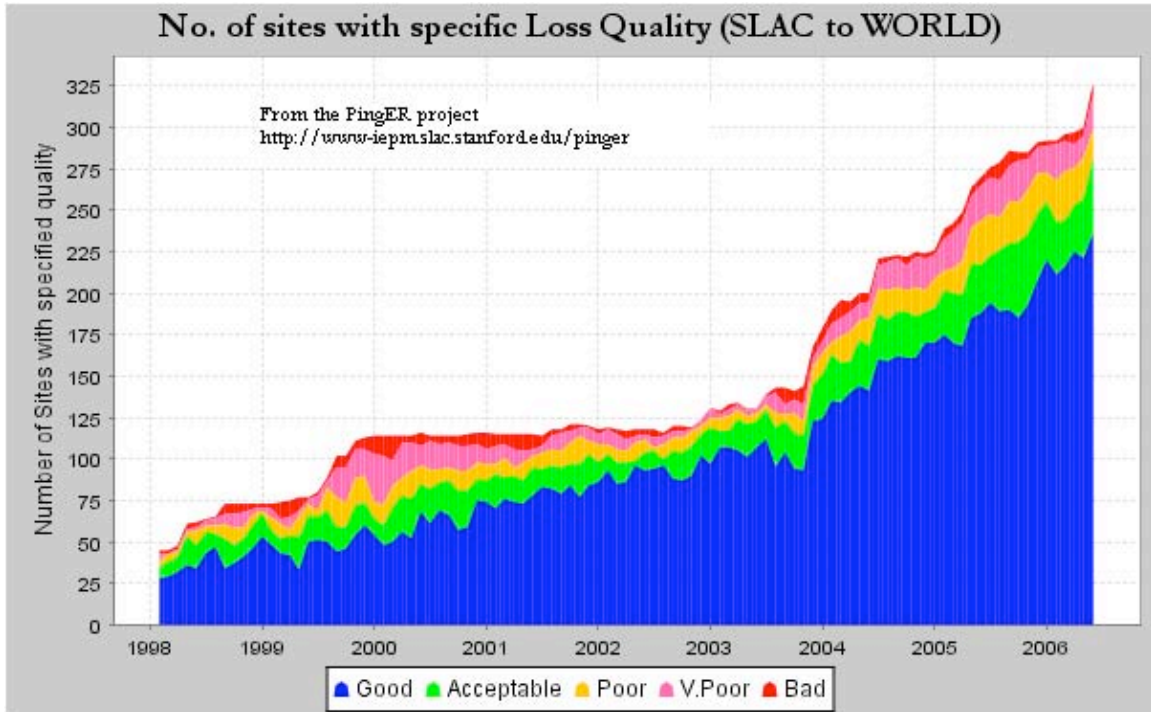


Figure 6: Number of hosts measured from SLAC for different loss qualities.

Another way at looking at this is to look at the losses by country and weigh this by the population in the country to yield the fractions of the world's population that are in countries with a given range of measured packet loss. This is shown in Fig. 7 for the turn of the millennium and December 2005. It can be seen that at the start of 2000, ~12% of the population lived in countries with acceptable or better packet loss. By December 2005 this had risen to 79%. The coverage of PingER has also increased from about 43 countries January 2000 to over 115 in December 2005. This in turn reduced the fraction of the connected population for which PingER has no measurements. The results are even more encouraging when one bears in mind that the newer countries being added typically are from regions that have traditionally poorer connectivity.

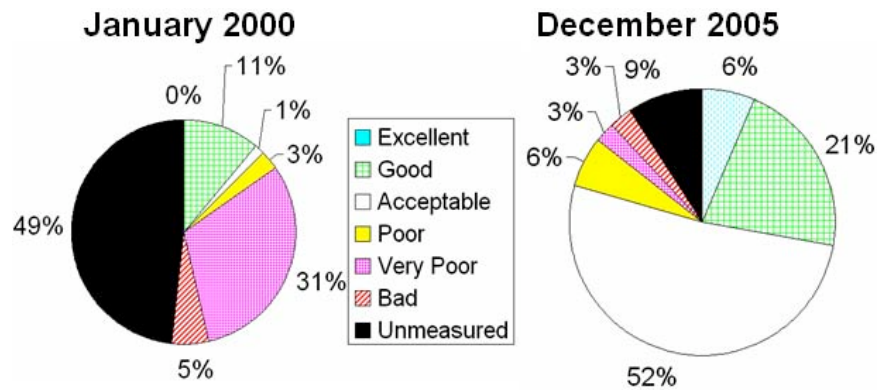


Figure 7: Fractions of the world's population in countries with given packet losses

Throughput Quality

Using the Mathis formula we can derive an estimate of the TCP throughput from our loss and RTT measurements. These estimates are only rough since the losses experienced by TCP²⁹ are different from those measured by ping, also PingER only sends about 14,400 pings a month between a monitoring host / remote host pair so one cannot see monthly losses of < 0.1% such as are often experienced on today's high quality paths. In addition the RTTs on high quality paths are approaching the limits of the speed of light in a fiber, so further improvement is difficult. None-the-less, especially for poorer quality paths, combining loss and RTT into a single metric is very useful. Fig. 8 shows the derived TCP throughputs measured from SLAC to the world's major regions, in some cases going back for the last 11 years. Similar plots (not shown here) are seen for the data measured from CERN in Geneva, Switzerland thus indicating that the effect is not just an anomaly associated with the measurements being from the U.S. The data for several of the developing countries only extends back for about five years and can vary greatly from month to month, so some care must be taken in interpreting the long term trends. With this caveat, it can be seen that links between the more developed regions including the U.S. and Canada, E. Asia and Europe are much better than elsewhere (3 - 10 times more throughput achievable). Regions such as Russia, S.E. Asia, S.E. Europe and Latin America appear to be 3-6 years behind. Russia and S.E. Asia are catching up slowly. However, Africa, S. Asia and C. Asia are 8-10 years behind and even worse appear to be falling further behind due to slow growth. Sites in many countries have less bandwidth than a residence in developed countries (typical residential DSL or cable bandwidths are of the order of a few hundred megabits/sec). Looking forward ten years to 2015, if the current rates of progress continue, then performance from N. America to Africa will be 1000 time worse than to Europe, to S. Asia and C. Asia will be 100 times worse than to Europe.

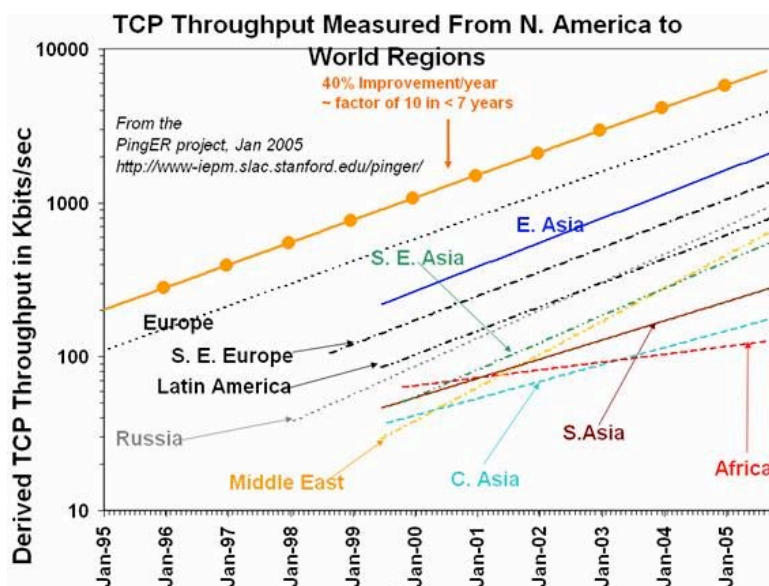


Figure 8: Derived TCP throughput as a function of time seen from the US to various regions of the world. The lines are exponential fits to the data.

²⁹ TCP deliberately provokes loss as part of its congestion detection algorithm.

Unreachability

The unreachability metric (a host is deemed unreachable if all ten pings in a set are lost) like loss is useful in that it identifies the fragility of the network paths or end hosts and is also largely distance independent (as opposed to say RTT or derived throughput). Fig. 9 illustrates the Unreachability for the regions of the world seen from SLAC in July 2006. As expected the developed regions of U.S., Canada, Europe, E. Asia and Oceania lead the way. Most regions are improving at a rate of approximately a factor of 10 in 8 years. The exception is Africa which is almost flat. The fragility of the networks in Africa and South Asia stands out, and the poor improvement seen for Africa is a cause for concern.

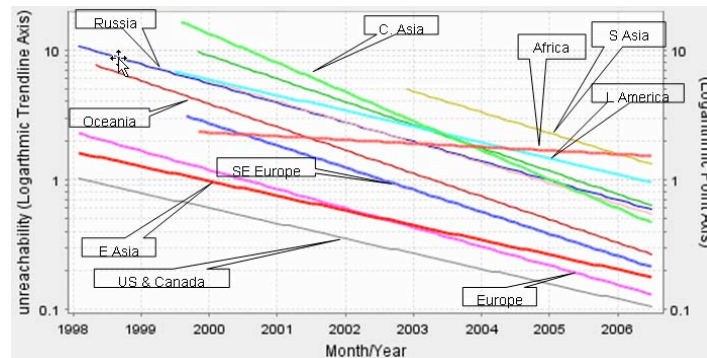


Figure 9: Unreachability for various regions as seen from SLAC.

Comparisons with Development Indices

Various human and technical development indicators have been developed by the U.N. and the International Telecommunications Union (ITU). It is interesting to see how well the PingER performance indicators correlate with the development indicators. The comparisons are particularly interesting in cases where the agreement is poor, and may point to some interesting anomalies or suspect data. Also given the difficulty of developing the human and technical indicators (at best they are updated once a year and usually much less frequently), having indicators such as PingER that are constantly and automatically updated is a useful complement.

One such Index that covers many countries and is reasonably up-to-date is the UNDP Human Development Indicator (HDI)⁴. This is a summary measure of human development). It measures the average achievements in a country in three basic dimensions of human development:

- A long and healthy life, as measured by life expectancy at birth
- Knowledge, as measured by the adult literacy rate (with two-thirds weight) and the combined primary, secondary and tertiary gross enrolment ratio (with one-third weight)
- A decent standard of living, as measured by GDP per capita.

Fig. 10 shows a scatter plot of the HDI versus the PingER Derived Throughput. Each point is colored according to the country's region. A logarithmic fit is also shown. Logarithmic is probably appropriate since Internet performance is increasing

exponentially in time and the differences between the countries can be related to number of years they are behind the developing countries, while human development is more linear. Since the PingER Derived TCP Throughput is linearly proportional to RTT, countries close to the U.S. (i.e. the U.S., Canada and Central America) may be expected to have elevated throughputs compared to their HDI. We thus do not plot or use these countries in the correlation fit between HDI and throughput. It is seen that there is a strong correlation ($R^2 > 0.6$) between the HDI and throughput. As expected countries in Africa generally occupy the lower values in x and y, and European countries together with Australia, New Zealand, Korea and Japan occupy the higher values of x and y.

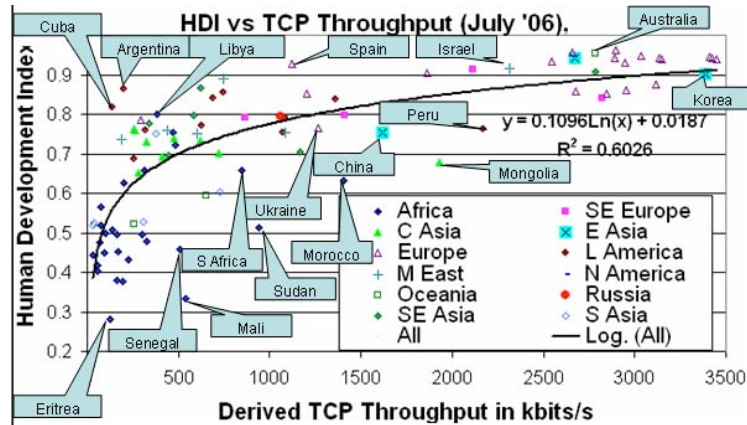


Figure 10: UNDP Human Development Index (HDI) vs. the PingER Derived TCP Throughput for July 2006

A second index introduced by the United Nations Development Programme (UNDP) is the Technology Achievement Index (TAI)³ that reflects a country's capacity to participate in the technological innovations of the network age. The TAI aims to capture how well a country is creating and diffusing technology and building a human skill base. It includes the following dimensions: Creation of technology (e.g. patents, royalty receipts); diffusion of recent innovations (Internet hosts/capita, high & medium tech exports as share of all exports); Diffusion of old innovations (log phones/capita, log of electric consumption/capita); Human skills (mean years of schooling, gross enrollment in tertiary level in science, math & engineering). The latest TAI is 2001, so it is older than the HDI. Also the TAI is only available for 50 countries that we have PingER data and that are not in North or Central America, whereas the HDI is available for 96 such countries. Fig. 11 shows a scatter plot of TAI vs. TCP throughput. In this case both the metrics are technological in nature so a linear fit works well. The correlation is seen to be strong and there are fewer outliers than in the HDI case.

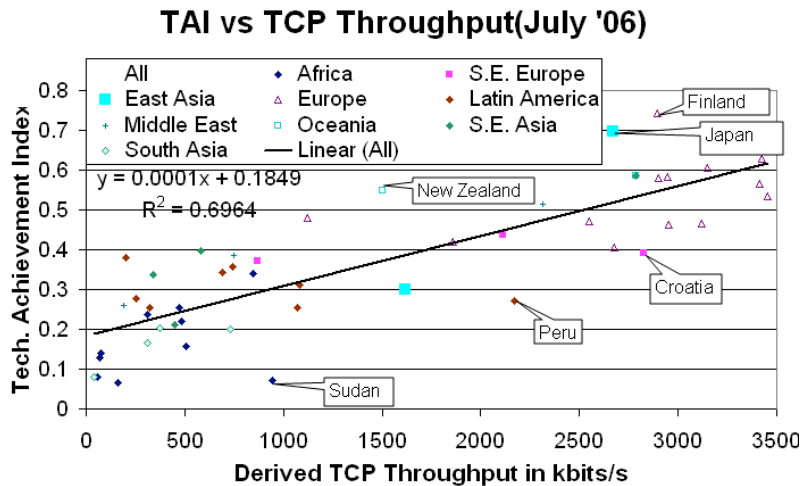


Figure 11: UNDP Technology Achievement Index (TAI) vs. the PingER Derived TCP Throughput for July 2006

Conclusions

The PingER project is arguably the most extensive network monitoring project active today in terms of deployment and longevity. It provides a simple, low network impact active end-to-end network monitoring methodology that is suitable for measuring a wide range of network performance. ViPER provides an easy to use interactive, appealing new user interface that can also be used in mobile clients. The PingER results illustrate the extent of the Digital Divide and how the gaps between the developed and developing countries are changing. The PingER results also complement other development indices and can be seen to correlate well.

Acknowledgement

We acknowledge all the people who have contributed to the development of PingER, in particular: Warren Matthews, Connie Logg, Aziz Rehmatullah, David Martin, and Maxim Grigoriev. We also acknowledge the assistance of the many administrators at the PingER monitoring sites who have helped install the toolkits, maintain the configurations, and respond patiently to requests for help to identify and fix problems. We are grateful for the support and funding from the U.S. Department of Energy, the U.S. State Department and the Pakistan Ministry of Science and Technology, as well as our parent institutions (SLAC and NIIT). Much inspiration for this work has come from: the International Center for Theoretical Physics (ICTP) and in particular Hilda Cerdeira, Enrique Canessa and Marco Zennaro; the International Committee on Future Accelerators (ICFA) Standing Committee on Inter-regional Connectivity (SCIC) and its chairman Harvey Newman; and the American Physical Society (APS) including its past chairman Burton Richter, and the chair elect of the APS Forum on International Physics, Herman Winick. Finally we thank Yee-Ting Li for careful reading of the paper and valuable suggestions.