2005 ANNUAL PROJECT REPORT

MAGGIE-NS

A COLLABORATIVE RESEARCH INITIATIVE BETWEEN
SLAC & NIIT

Duration:

September 2004 – August 2005

Project title:

Measurement and Analysis of the Global Grid and Internet End-to-End Performance-NIIT SLAC (MAGGIE-NS)

Principal Investigators:

Dr. Roger L. A. Cottrell
Assistant Director
SLAC Computing Services
Stanford Linear Accelerator Center (SLAC),
USA
cottrell@slac.stanford.edu

Prof. Dr. Arshad Ali
Director General
NUST Institute of Information Technology (NIIT),
Pakistan
arshad.ali@niit.edu.pk

Project Website:
http://maggie.niit.edu.pk/

Table of Contents

Executive Summary

The growth in Distributed and Network Systems provokes the challenge of developing a lightweight and reliable monitoring infrastructure, appropriate for the bandwidth constrained environments like Pakistan. MAGGIE-NS steps up to the task of developing and deploying solutions related to Network Monitoring, Network Troubleshooting, Capacity Planning and Auditing, Grid Scheduling and Evaluation of experimental networking components. MAGGIE-NS primarily aims at deploying a suitable monitoring infrastructure at NIIT and hosts in other developing countries, which is robust enough to be comparable to other monitoring infrastructures around the world, yet lightweight enough to generate minimal traffic on the already congested links. In order to solve these problems, MAGGIE-NS project has made extensive improvements, and extended the deployment and use of two major network monitoring infrastructures developed by the SLAC Internet End-to-end Performance Monitoring (IEPM) team: PingER and IEPM-BW.

PingER was installed as a monitoring host at NIIT in Pakistan, and other developing countries including South Africa, Brazil and India, to increase the global coverage; the number of monitoring sites was increased by 5%, monitored sites by 20% and countries monitored by 10%. Due to the nature of the collaboration, a prime focus was on the network infrastructure of Pakistan hence a dozen remote sites in Pakistan were also included in the NIIT monitoring list. This enabled NIIT to measure its performance to geographically important locations in the world as well as get a more detailed overview of the connectivity with various parts of the country. In order to achieve sustainable growth, the installation procedure of PingER had to be simplified and automated PingER management modules were developed for filtering the unreliable nodes and data. Moreover, visualization of the global network monitoring information provided by PingER was significantly improved with a view to present the graphs to the higher management for the purpose of decision making. Currently, efforts are being made to integrate the display of the PingER data with MonALISA to enable more close to real-time display of the data by a monitoring node.

IEPM-BW was developed/deployed to provide active End-to-End (E2E) application and network measurements to assist with detailed diagnostics about the network, gathering parameters like link bandwidth (available and utilized) and routing information etc. The MAGGIE-NS project facilitated the installation of IEPM-BW at NIIT; making NIIT the only IEPM-BW node outside Europe and North America. IEPM-BW and PingER together help in understanding the routing and congestion problems in the country's network.

In this report, the various independent student projects are also discussed which have contributed in identifying network anomalies and either support or make use of data provided by IEPM-BW and PingER.

TULIP (IP Locator Using Triangulation) is a project to generate pings from various landmark sites around the world to an unknown IP location, in order to correctly determine the exact geographical location of that particular IP, by estimating the Great Circle Distances (GCDs) derived from minimum Round Trip Time (RTT) measurements.

Another tool, traceanal has been developed which helps in detecting and visualizing the route changes. This is a fairly common problem in the regions where the network infrastructure is poor as is the case in most countries on the unfortunate side of the digital divide.

Network Weather Forecasting is used to predict the behavior of the network by extrapolating various parameters. It enables the network administrators and stakeholders to visualize the expected behavior of the network based on its history. Implementation of various forecasting and anomaly detection algorithms has been carried out to detect significant, persistent anomalous events in real E2E Internet Performance measurements. Besides active probing, methods for passive monitoring and anomaly detection are also being investigated in order to ensure addition of minimal traffic on the internet. Progress has been made on detecting network anomalies using the Principal Component Analysis. The results of this analysis are very encouraging especially given its ability to simultaneously look at multiple paths and multiple metrics.

Connectivity for the Pakistan side of the project suffered during the Fiber Outage in Pakistan from the June 27th 2005 to July 8th 2005. However, the MAGGIE-NS team at SLAC was quick to grab this opportunity to study and analyze the effect of this loss in terms of overall network performance of the country. This also provided an opportunity to study the overall network infrastructure of Pakistan in general, and NTC and NIIT in particular. We also identified and reported that the overall quality of service experienced by Pakistani universities is poor. This analysis can be extremely helpful for the Pakistani Higher Education Commission (HEC), as well as Pakistan's National Telecommunication Corporation (NTC) as the latter is the service provider for the Pakistan Education and Research Network (PERN).

MAGGIE-NS project opened the doors for the research community at SLAC to identify and understand the issues in the network infrastructure of Pakistan. However, the contributions are much beyond the bits and bytes. The culture of research that has been established at NIIT, with SLAC playing its due part, has been phenomenal in the overall growth of the institution. The cross visits by the researchers at SLAC and students and faculty at NIIT has proved extraordinary in enhancing the vision and expertise of these individuals. These experiences will have a very beneficial impact on understanding and improving the performance of Pakistani networks, proving understanding for support people and researchers, help sparking an Information and Communication Technology (ICT) revolution in the country, and benefit Pakistani society by leveling the connectivity playing field and improving access to knowledge.

MAGGIE-NS project has resulted in publication of two research papers, one report and twelve talks and presentations ranging from Network Problems to the quantification of digital divide. It is hoped that the extension of MAGGIE-NS and a more focused proposal on the network infrastructure for PERN, called eMAP (end-to-end Measurement Architecture for PERN) will be funded so that the promising work initiated can be fully utilized for the benefit of Pakistan as well as the United States of America.

1. Introduction

Scientific research has greater computing and networking power than ever before, but it is recognized in the scientific community that the increases in power also bring with them major challenges. The deployment of Grid technologies and a production quality Grid Service requires detailed information on network performance. It is widely believed that what is needed is to develop a ubiquitous monitoring infrastructure that would not only provide access to the measurements seen in multiple monitoring projects but also provide the novel addition of allowing us to co-ordinate and integrate tools in a co-operative framework. Such a requirement is particularly needed in Pakistan since it is initiating and collaborating in a number of Grid computing projects like the National University of Sciences and Technology (NUST[1]) and Quaid-e-Azam University[2] involvement in the CERN Large Hadron Collider (LHC) Computing Grid (LCG) collaboration. In addition, proposals have been submitted to develop an e-education and distance learning environment in Pakistan. Looking at the growth in the Grid Technologies, MAGGIE-NS proposed to integrate numerous network and application performance monitoring tools into a scalable infrastructure providing measurements, analysis and access to data. This enables us to gather useful data from different network monitoring and measurement tools in a closer to real-time manner for analysis. This data gathering could be utilized to address the above mentioned Grid needs as well as to: quickly see long term-trends; compare the performance of different world regions; spot problems such as sudden, significant, unexpected changes in performance (anomalies); set expectations (for example do Pakistan Research and Education sites have adequate network infrastructure to participate in LHC Computing Grid activities); and provide higher level decision making information for upper level management.

2. Objectives & Scope

The major objectives of the Maggie-NS project were to propose and deploy solutions related to Network Monitoring, Network Troubleshooting, Capacity Planning and Auditing, Grid Scheduling and Evaluation of experimental networking components.

The objective of network troubleshooting is to quickly detect and report troubles and to provide information as to the magnitude of the trouble, when and where it occurred and details to assist in isolation and by-pass. This requires continuous monitoring, archiving and analysis of the data to determine normal behavior and to identify and locate anomalies in time and space.

Network monitoring data is necessary for capacity planning, so that resources are readily available when needed, and in order to prevent bottlenecks. Monitoring data is also needed to set realistic performance expectations and to define measurable service level agreements that can be agreed upon by the providers and users. Once the service level agreements are in place, monitoring is required to audit and ensure that the service level agreements are being met.

Nodes in Grid Computing are spread across the globe and scheduling the jobs on different Grid nodes requires comprehensive monitored data; for instance in order to

provide optimal network access to data repositories and other resources, decision making based on the network monitoring data. The purpose of the project was to provide the data in a format which could be accessed by the Grid Services and used for decision making in scheduling the Grid related jobs.

Evaluation of networking components is an important factor of network reliability and performance. So the network monitoring data being collected could be used for evaluation of: new protocols for very high-speed and low-speed networks, new network interface card (NIC) features, large maximum transfer units (MTUs), optimal configuration parameters (such as window sizes and number of parallel streams) etc.

3. Network Monitoring

Current Grid applications typically use a much smaller percentage of available network bandwidth than expected. Whereas the application developers often blame the network as a problem for poor application performance, network engineers typically point to host issues or poor application design. Network monitoring services are necessary to identify and resolve the network issues and verify whether the network is in fact the source of the problem. Discovering and fixing Grid application configuration and performance problems is a challenging task. Problems can manifest themselves in many shapes and forms; making it difficult even for the experienced network engineers and application developers to understand the extent if such problems.

The MAGGIE-NS project has made extensive improvements, and extended the deployment and use of two major network monitoring infrastructures to assist in tackling the above challenges, The infrastructures are PingER and IEPM-BW.

3.1. The PingER Project

As part of MAGGIE-NS, PingER[3] was installed as a monitoring host in Pakistan (at NIIT), India, South Africa, and Brazil. In addition the coverage of remote sites monitored was extended by about 20% to 673 and the countries monitored by about 10% to 114. The aim was to improve the overall understanding of the network performance within and from Digital Divide regions and in particular, Pakistan.

Since PingER measurements[4] add very little network traffic (< 100bits/s on average per monitor-remote host pair) it can be very effectively used for monitoring low-performance links such as those available to, within and between developing countries and regions such as Pakistan, India, Russia, Latin America, and Africa where the network infrastructure is less advanced. With the installation of PingER monitoring sites in Pakistan, India, Brazil, Russia and South Africa, the PingER/MAGGIE-NS project has made significant advances on quantifying the Digital Divide in terms of Internet performance[5]. The addition of about a dozen monitored remote sites in Pakistan has also enabled a more detailed study of the connectivity within Pakistan.

3.1.1 PingER Management

Since its inception, the size of the PingER project has grown to where it is now monitored hosts in over 110 countries from about 35 monitoring hosts in 14 countries. With growth in the number of monitoring as well as monitored (remote) nodes, it was

perceived that automated mechanisms need to be developed for managing this project. The following modules for PingER management project are being developed or under testing at NIIT/SLAC as part of the MAGGIE-NS project:

- Creation of filters to indicate the monitoring sites whose data is not available
- Creation of filters to indicate the monitored sites that are not available and categorize them according to their response status.
- Identification of a host that physically moves to a new location (e.g. a named web server actually is a proxy that is not where it used to be).
- Automated report generation tool to generate daily, monthly, yearly reports regarding problems in monitored data.
- Detect sudden, significant (anomalous) changes in the behavior (including breaks in reachability) of the network.
- Identifying discrepancies (e.g. impossible values) in measured data and in the host configuration databases (e.g.; at the time of registration of the monitored hosts, the data entered might be incorrect and incomplete).

3.1.2 PingER Installation

Before the start of MAGGIE-NS project, PingER had a complex installation procedure. An initial improved installation process was developed by students working with Warren Matthews at Georgia Tech. This was extended, and productized by two MAGGIE-NS project students in order to integrate the improvements and make PingER easier to install for the monitoring sites. This upgrade was necessary, given the increase in the number of monitoring sites around the globe, and the lack of technical skills at the newer sites. The new version is called PingER2, which possesses the same functionality as PingER, but is much easier to install. As a result, nodes in Africa and Turkey are installing PingER2, which will be a great addition to the coverage. Being an important collaborator, NIIT is monitoring all the newly added nodes. The new results will pave way for an interesting study of the digital divide from the countries and regions that are on the unfortunate side of the digital divide. Essentially, this will help in classifying the "divide that exists within the digital divide", i.e. the division between which developing countries/regions having the worst performance.

3.1.3 PingER Visualization

To assist in visualizing and accessing PingER data, the MAGGIE-NS project has improved the storage of configuration data by converting to a MySQL database and adding several new fields of information per host. This in turn has enabled us to develop mouse sensitive maps that graphically illustrate the deployment of PingER.

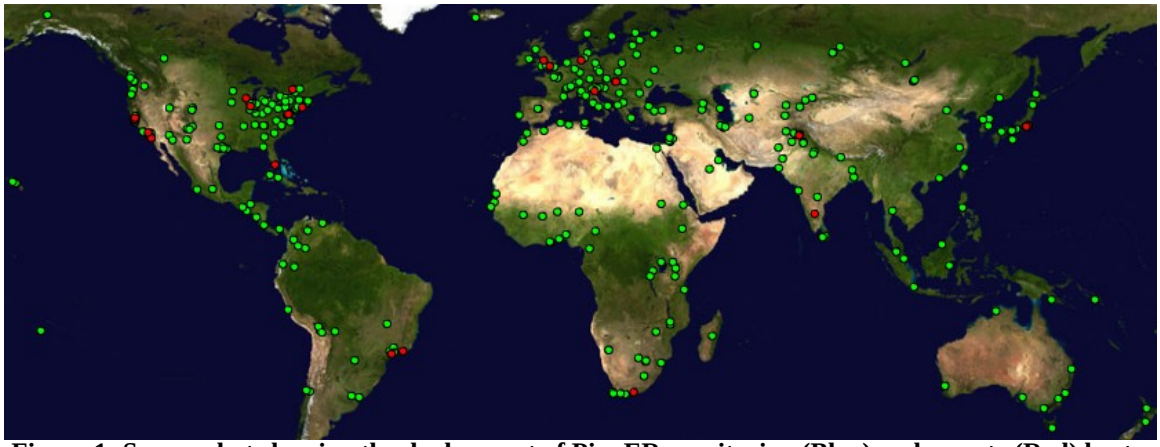

Figure 1: Screen shot showing the deployment of PingER monitoring (Blue) and remote (Red) hosts. When used interactively this map allows one to move the mouse over the host to find the name of the host.

MAGGIE-NS has developed several scripts to assist in the creation of executive level figures showing aggregations of performance by regions, correlations vs. populations etc. Prior to developing these scripts, creating such figures was very labor intensive. The MAGGIE-NS project is currently working on integrating the display of the PingER data with MonALISA[6] to enable more close to real-time display of the data by a monitoring node. This will be a big improvement over the current simple tabular display functionality for real time information.

3.2 The IEPM-BW Project

The primary task of the IEPM-BW[7] was to develop/deploy a simple and robust, ssh based active End-to-End (E2E) application and network measurement and management infrastructure. The IEPM project that was funded by the DOE/MICS base program has the goal of providing Internet E2E performance monitoring for planning, setting expectations, problem identification trouble-shooting, and networking research. The MAGGIE-NS project has deployed IEPM-BW at BNL, CERN, FNAL, SLAC and NIIT. The first four are aimed at the High Energy Physics (HEP) CERN/LHC, FNAL/CDF/D0 and the SLAC/BaBar experiments. The NIIT installation is important not only for understanding the connectivity of NIIT and Pakistan with the rest of the world, but also since it is the only IEPM-BW monitoring node outside Europe and the US.

3.2.1 Importance of IEPM data

PathChirp: As part of the MAGGIE-NS project, IEPM-BW has been extended to support PathChirp[8]. This is a packet pair dispersion technique for measuring available bandwidth. Due to its improved accuracy[9] compared to ABwE[10], focus is now on PathChirp to dynamically estimate the available bandwidth along an E2E network path from several monitoring sites to other remote sites.

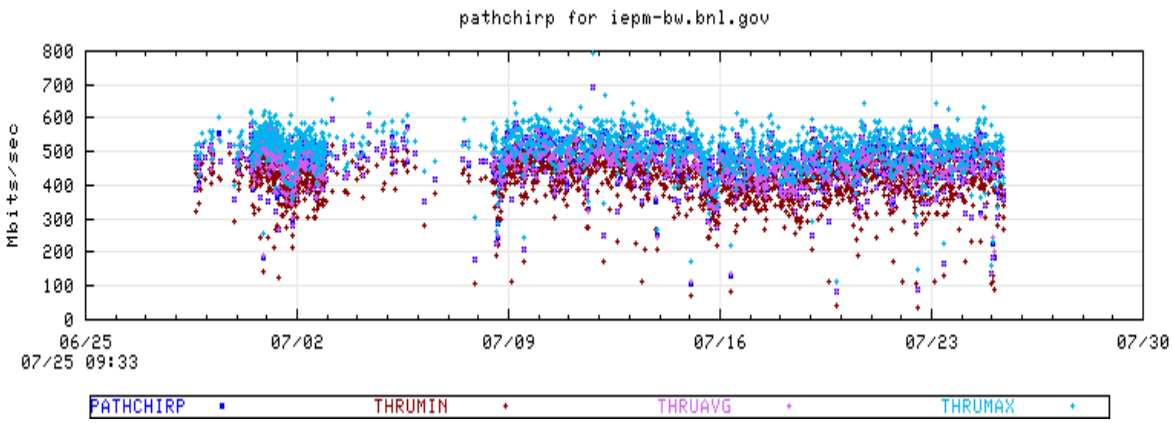

Figure 2 Example of pathchirp results, for minimum, average and maximum available bandwidth (thru) from iepm-bw.slac.stanford.edu to iepm-be.bnl.gov

Traceroute: Traceroute provides a hop by hop list of the routers along the path and the RTT to each router. Visualization of traceroute is described in section 5.2.

Ping: The main information provided by ping is the reachability, packet loss and RTT. RTT information is obtained in IEPM-BW using Ping measurements. An example of ping RTT time series is shown in Figure 3. The measurements are made at regular 10 minute intervals. Missing points indicate that the remote host (NIIT) was unreachable. The minimum RTT (blue dots) is around 500ms. It can also be seen that at times congestion occurs during the whole sample of 10 pings measured so the minimum RTT of the sample varies with time (increasing during work hours).

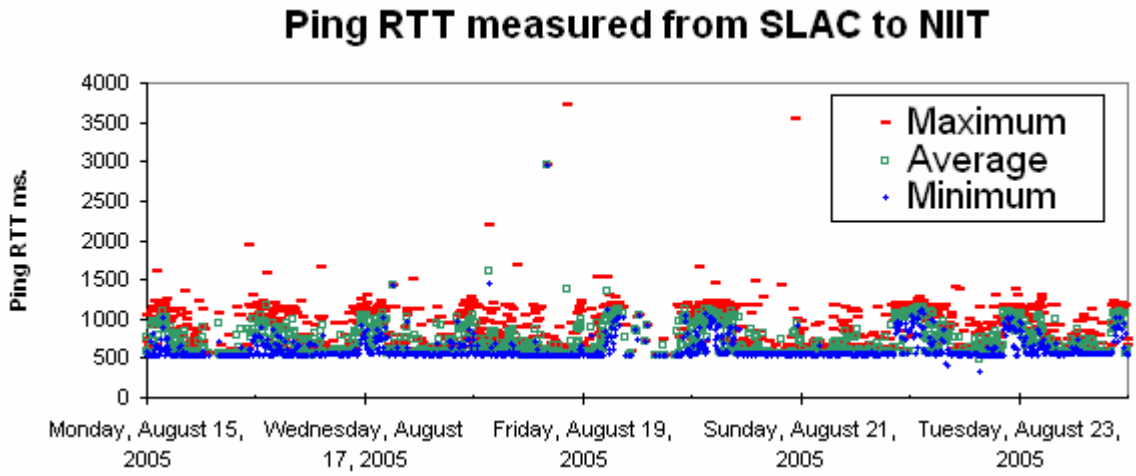

Figure 3 Example of the ping minimum RTT, average RTT and maximum RTT from iepm-bw.slac.stanford.edu to node1.niit.pk

Iperf: Iperf[11] is a tool that used to measure TCP achievable throughput. It allows the tuning of various parameters such as window size and number of parallel streams. In IEPM-BW h iperf with a single stream and with multiple streams (miperf) are run to measure TCP achievable throughput.

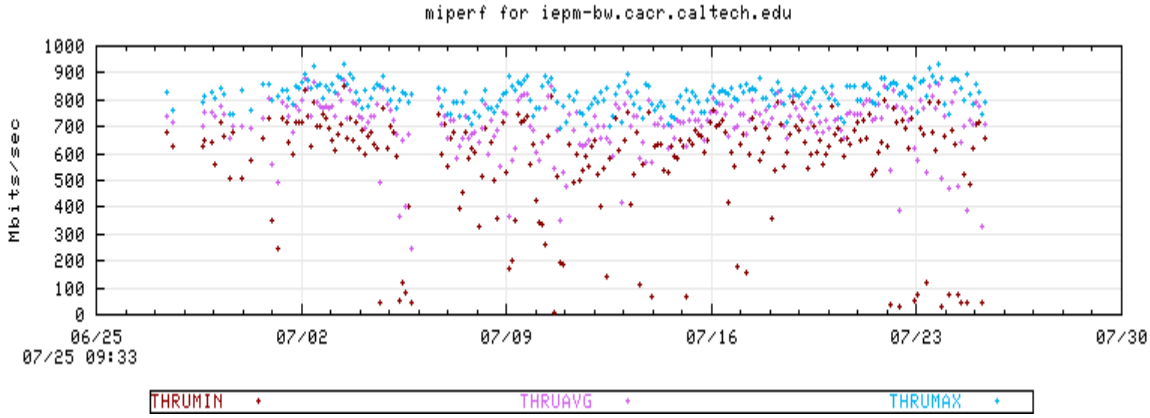

Figure 4 Multi stream iperf minimum (flow) achievable TCP throughput, maximum throughput and average throughput from iepm-bw.slac.stanford.edu to iepm-bw.cacr.caltech.edu

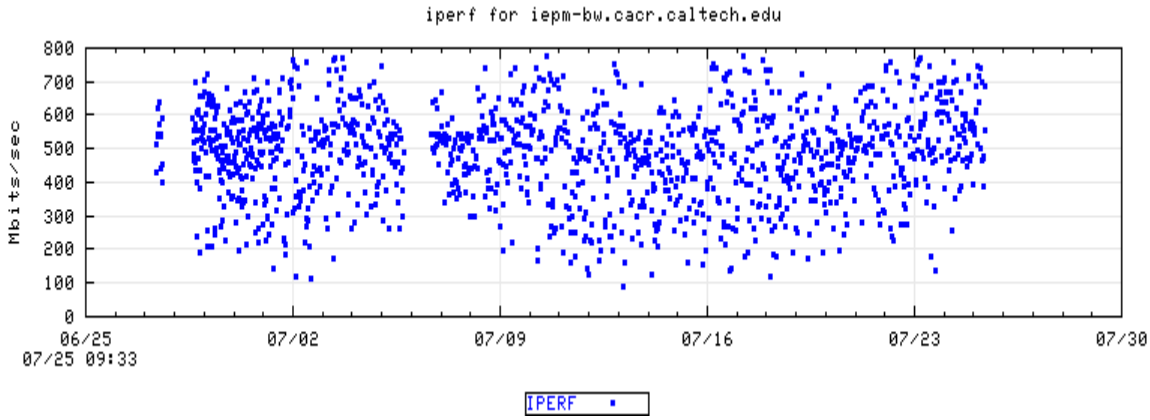

Figure 5 Iperf (single-stream achievable TCP throughput) measurement from iepm-bw.slac.stanford.edu to iepm-bw.cacr.caltech.edu

Thrulay TCP: As part of the MAGGIE-NS project IEPM-BW has been extended to support thrulay[12] TCP achievable throughput and RTT measurement tool. It is used to measure the capacity of a network by sending a bulk TCP stream over the path. Thrulay can report the TCP achievable throughput and RTT at intervals (e.g. 1 second) during its measurement duration.

In Figure 6, the measurement durations were 15 seconds and the rttmin etc. are the minimum, average and maximum of the 15 one second intervals at which RTT was reported. It shows the achievable throughput from SLAC to CERN.

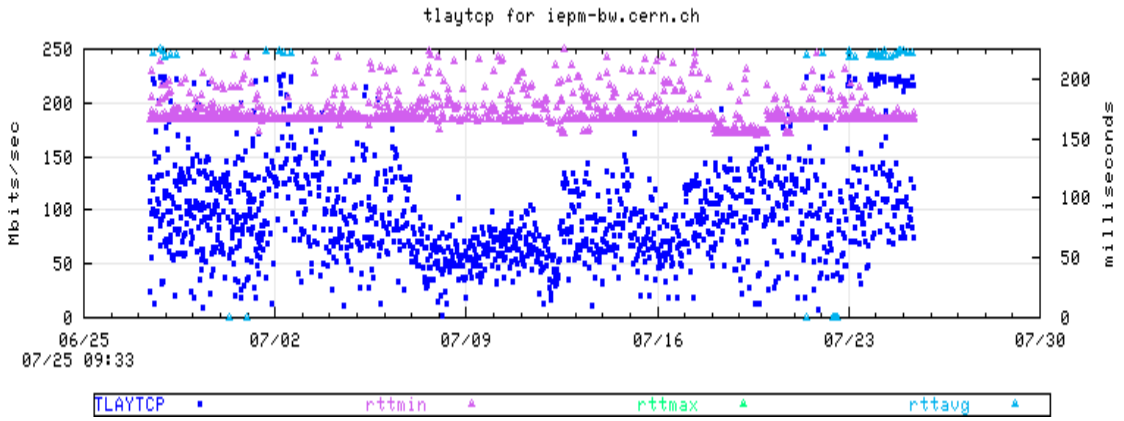

Figure 6 Measured achievable throughput as tlaytcp, together with rttmin, rttmax and rttavg from iepm-bw.slac.stanford.edu to iepm-bw.cern.ch

3.2.2 IEPM Deployment

In addition to the 5 IEPM-BW monitoring sites, there are 40 remote sites in 10 countries. Figure 7 shows the monitoring site for SLAC and the remote IEPM-BW sites it monitors around the world. This figure was developed as part of the MAGGIE-NS project by NIIT students in an attempt to make IEPM-BW more user-friendly and make analysis easier by providing an advanced visualization tool. Prior to the deployment of this tool, IEPM-BW real time data was available only in text format. It is also an example of a real-time visualization (lines shown illustrate the current measurement path; mouse-over displays the most recent results). At the time of snapshot, ping and thrulay tests were in progress as depicted by the yellow and green lines. This tool will also be used as part of a demonstration at SC2005[13].

Figure 7 Deployment of the IEPM-BW hosts monitored from SLAC. The monitoring site (SLAC) is shown in blue, and the monitored sites in red. The colored lines indicates the test being made when the screen snapshot was taken

3.2.3 IEPM-BW Data Visualization

IEPM-BW can measure numerous parameters from various tools like ping, traceroute, pathchirp, thrulay, iperf, miperf etc. A considerable fraction of the work in MAGGIE_NS has been focused on making these measurements easier to visualize. In addition to the individual metric, time series plots for the last 30 days have also been implemented, IEPM-BW also provides plots of time series for multiple metrics, longer term historical (6 weeks and 6 months) time series, frequency histograms of the metric distributions, scatter plots comparing various metrics, traceroute visualization, tabular details

concerning the latest measurements, as well as access to the raw data to enable further analysis. In addition there is extensive visualization of the IEPM_BW management information, problems encountered etc.

4. Infrastructure management to enable sustainability

The Maggie-NS project is focused on studying the network infrastructure of Pakistan in general and NIIT in particular. Therefore, PingER has been installed as a monitoring site at NIIT, and the data gathered is being studied comprehensively to understand the various connectivity patterns in Pakistan. NIIT is acquiring its Internet connectivity services from NTC[14], which is also the official Information Technology and Telecommunications service provider to the government of Pakistan and its related organizations. Moreover, NTC is also servicing PERN[15] which connects most educational and research sites in Pakistan[16], NIIT connectivity is thus a representative of the connectivity of a number of major sites in the country.

Monitoring the major educational institutes in Pakistan from NIIT was a focal point of MAGGIE-NS. One interesting observation was that in the last year, several nodes (in particular the major web servers for many sites) have shifted from within Pakistan to outside of the country. Two examples are HEC[17] and Bahria University[18], Islamabad. This was discovered by looking at the minimum ping RTT times to these sites from SLAC and confirmed by looking at the traceroutes. The minimum measured RTT values were smaller than possible knowing the distance and the speed of light in a fiber.

5. Major Projects Addressed by MAGGIE-NS

5.1 TULIP- IP Locator Using Triangulation

With the extensive growth in the coverage of PingER arises the great difficulty of keeping track of the changes in the physical locations of the monitored sites. This might lead to mis-leading conclusions. To assist in automated location verification of PingER hosts, MAGGIE-NS has launched a task to build a tool to give the latitude and longitude for a given IP address or URL. This tool will then be used to identify hosts whose located position is in conflict with the PingER database latitude and longitude by comparing the minimum RTT with that predicted from distance between the monitor and remote sites. Anomalies will be reported so the PingER database can be corrected using values from the locator tool and/or new hosts can be chosen to be monitored.

The location of an IP address can be determined using by using the minimum RTT measured from multiple “landmark” sites at known locations, and triangulating the results to obtain an approximate location. The basic tool will be a script that takes RTT measurements from landmarks to a selected target host (typically at an unknown location) specified by the user and figures out the latitudes and longitudes of the target host.

TULIP is a project to implement this triangulation technique. TULIP requests the landmark sites to make RTT measurements, collects these measurements, analyzes the data, presents the results and on demand makes traceroutes to the remote host from elected landmark sites to enable the user to see the paths. .It is under development and will also utilize the historical min-RTT PingER data in order to verify the locations of

hosts/sites recorded in the PingER configuration database, and to optimize the choices of parameters used by TULIP.

For landmarks we will: use existing LookingGlass sites (that allow on demand ping measurements); add ping web agents to PingER monitoring sites and are working with the authors of [19] to utilize their landmark sites.

While implementing such a technique there are several factors that have to be taken into account:

- **Bandwidth:** Links can have capacity bandwidths that differ by many orders of magnitude (e.g. 1 Mbits/s versus 10Gbits/s). For a 1 Mbits/s link the time to clock a 100byte packet onto the link is about 0.8msec. This is roughly the delay time for light to propagate through a 100 mile (~161km) fiber (or roughly[20] 100km per msec of RTT, n.b. RTT ~ delay * 2). Thus, It may not be possible to have a single conversion factor (alpha) from RTT to distance that works for all links for all distances.
- **Queuing delay** can have major effects on the RTT measurement. Queuing delays are critically affected by **cross-traffic**. Increased cross-traffic is likely to increase queuing and thus RTTs. Hence the minimum RTT of multiple measurements is used with the assumption that the minimum possible RTT is for packets that see no queuing. However if the queuing persists for the duration of the measurement then the minimum RTT will include a queuing component and not accurately reflect the distance.
- **Routing policy:** The path a packet takes over the network may change dependent on link/router availability and routing policies. Extra hops in the path will add extra router and clocking delays, different hops will have different delays, and different links will have different capacities and distances. All of this will change the minimum RTT when a route changes. Further the peering may be such that the route between a landmark and a selected host is not direct. For example, the route from two landmarks in San Jose and Palo Alto to SLAC (< 20 miles away) goes via New York. Thus one may need to compare the results from a cluster of landmarks that are close together (but have diverse peering) and choose the minimum RTT.
- **Bottleneck bandwidth** is the available bandwidth of the link in the path that at the measurement moment has the lowest available bandwidth. This may be the link with the lowest static capacity or it may be a heavily loaded link with the currently minimum available bandwidth (~ link capacity – link utilization). The link utilization changes as a function of time so the bottleneck magnitude may vary and its location may move from link to link.
- **Misconfiguration of network devices** such as routers and switches can cause unexpected delays. First of all misconfigurations can cause a router or link to fail. Secondly the alternative route may be poorly chosen (e.g. have much lower capacity or much longer delays).
- **Traffic conditions** may vary enormously as a function of time. For example, on Monday mornings typically there is a surge in traffic as people come to work; also

there is usually much more traffic during peak work hours compared to weekends and nights. Such increased traffic often leads to congestion, increased queuing and thus increases in RTT.

Links in Canada, USA, Europe, Japan are generally of high capacity e.g. 100Mbits/s or 1 Gbits/s end host Network interfaces, 1 or 10Gbits/s backbone links and >= 100Mbits/s border router links. Also the routes tend to be well defined with optimum backup paths. Thus the distances derived from the minimum RTT values in those regions are generally accurate.

However, many developing regions/countries such as Pakistan do not yet have a similar infrastructure. For example some paths between sites in Pakistan have been observed by the MAGGIE-NS project to go via Europe or E. Asia. Furthermore, as mentioned above, the assumption that 1ms of RTT for a distance of 100km on the fiber may not work well for short paths with slow links.

Even for well developed infrastructures, the actual RTT is usually larger than that given by 100km = 1 ms, since between nodes the path is often not a straight line (often following roads, railways and other rights of way), and the nodes themselves are not usually on a great circle route. Further there are delays in the routers themselves. A more realistic starting value is therefore for the conversion factor (alpha) to be closer to 50 km = 1 ms. One of the goals of this project will be to optimize alpha to optimize the agreement of RTT distance estimations versus the GCDs between the sites and to understand whether multiple values are needed for alpha.

Table 1 below shows the minimum RTTs values from some of the Looking Glass servers[21] landmarks to www.slac.stanford.edu (located at SLAC in Menlo Park about 40 miles south of San Francisco) using an alpha of 50 km / ms. It also illustrates some of the challenges of this method.

- As mentioned above, for certain services providers, the path between the SLAC and the landmarks in Palo Alto and San Jose goes via New York.
- Burnaby is located a few miles from Vancouver, Canada. Pings from SLAC to the University of British Columbia in Vancouver have a minimum RTT of about 24ms, equivalent to a distance of about 1200km in better agreement with the GCD. The route to Burnaby goes via Seattle (20 ms) and then the next hop extends to 89ms. Although the cause of this delay is being investigated, such problems are a common-place in real networks.
- The minimum RTT for Chicago results in a distance that appears to be too large. The traceroute shows the path from SLAC to this host goes via Avenue of the Americas in New York, hence the extra delay.
- The minimum RTTs to Singapore, Malaysia, and the Philippines predict shorter distances than the GCD. This is even though the routes go via Chicago which will add an extra 60 ms. This indicates that the trans-Pacific routes are closer to great circles as compared to the paths across land (where the path follows rights of way etc. and the routers are not usually on the great circle route). In this case the value for alpha of 50km is probably too small.

Ignoring the landmarks in Palo Alto, San Jose, Burnaby and Chicago, and using alpha = 50 km, the average difference = (GCD – min_RTT_distance) / GCD ~ 10% ± 9%. If an alpha of 55km is used, the average difference would be ~ 0.

Table 1 Minimum RTTs for one way from selected landmarks to www.slac.stanford.edu

Serial No.	Landmark location	Min RTT ms	Avg RTT ms	Max RTT ms	Great Circle Distance (GCD) km	Estimated Distance from Min RTT km	(GCD - RTT dist) /GCD
0	PaloAlto	75	84	100	4	3752	
1	SanJose	74	83	98	4916	3711	
2	Chicago	80	80	81	2981	3975	-33.3%
3	Washington DC	74	74.2	75	3922	3699	5.7%
4	NewYork	79.5	79.6	79.7	4126	3977	3.6%
5	Stanford	0	null	null	3		
6	Oslo	0	null	null	8332		
7	Israel	216	219	228	11931	10800	9.5%
8	Geneva	164	164	168	9394	8200	12.7%
9	Amsterdam	160	164	168	8768	8000	8.8%
10	Switzerland	152	153	156	9365	7600	18.8%
11	Toronto	88	92	96	3656	4400	-20.4%
12	Burnaby	88	92	104	1276	4400	
13	Dublin	152	164	208	8172	7600	7.0%
14	Milan	168	169	172	9564	8400	12.2%
15	Brussels	152	156	176	8872	7600	14.3%
16	Budapest	176	177	180	9786	8800	10.1%
17	Cologne	0	null	null	8980		
18	Copenhagen	164	167	172	8779	8200	6.6%
19	Dusseldorf	164	168	172	8948	8200	8.4%
20	Frankfurt	160	162	164	9126	8000	12.3%
21	London	144	152	172	8609	7200	16.4%
22	Moscow	196	208	220	9439	9800	-3.8%
23	Oslo	172	176	180	8332	8600	-3.2%
24	Paris	148	148	152	8945	7400	17.3%
25	Prague	172	172	172	9364	8600	8.2%
26	Vienna	172	177	192	9611	8600	10.5%
27	Zurich	172	174	176	9365	8600	8.2%
28	Munich	0	null	null	9428		
29	Bucharest	188	190	192	10335	9400	9.0%
30	Paris	152	152	152	8945	7596	15.1%
31	Torino	167	167	168	9529	8335	12.5%
32	Singapore	184	193	212	13577	9200	32.2%
33	Malaysia	212	217	228	13581	10600	21.9%
34	Thailand	220	220	224	12739	11000	13.7%
35	India	236	239	240	12341	11800	4.4%
36	Philippines	172	174	176	11207	8600	23.3%
37	Pakistan	0	null	null	11902		

More landmarks are being added and a prototype of the application is almost complete. More paths using the known locations of PingER monitoring and remote sites and the measured minimum RTTs are being investigated. The GeoLIM Constraint-Based Geolocation technique[19] will also be used to optimize the localization and assign confidence regions using measurements from multiple landmarks. Especially with the increase in number of nodes, this utility will positively contribute towards the automation of selecting appropriate PingER hosts. Moreover, policies have to be defined for the variation of Alpha based on many factors.

The parallel GeoLIM project has deployed their own RTT measurement agent at landmarks in Western Europe and the U.S. The limited distribution of GeoLIM landmarks has resulted in poor localization for hosts in developing regions. Hence TULIP would focus upon providing measurements from a much wider distribution of landmarks (including Looking Glass servers, PingER monitoring hosts etc.) Moreover, efforts are also being made with the GeoLIM project people to assist in deploying their agent at PingER sites and sharing their landmark agents.

5.2 Traceroute Change Analysis (TCA)

Route change analysis is critically important for network monitoring and troubleshooting. A route change can be a major reason for fluctuations in round trip times, achievable and available bandwidth, file transfer throughput etc. However, an important point to consider is that not all route changes cause throughput changes. An example of a throughput change correlated with route changes, can be seen in Figure 9.

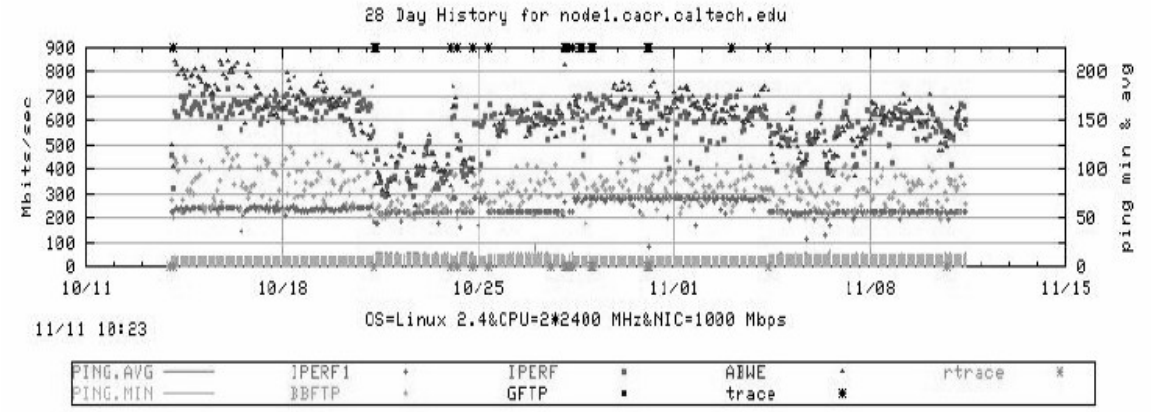

Figure 8 Time series plot with route changes indicated

In Figure 8, the asterisks along the top axis indicate forward path traceroute changes. The asterisks along the bottom axis indicate reverse route changes. Note the correspondence between throughput change on 10/21/05 and the forward route change.

As part of MAGGIE-NS, a new tool traceanal[22] has been implemented which is used to detect and visualize the route changes, and publish the results of route analysis through a web link[23].

The web accessible route daily summary page (Figure 10) is created and updated throughout the day. At the top of the page, there are links to "Yesterday's Summary", today's "Reverse Traceroute Summary", and the directory containing the historical traceroute summaries. Under those links is the traceroute summary table which provides "at a glance" visualization of traceroute change patterns. This facilitates the observation of synchronized route changes for multiple hosts in the cases that a common sub route changes for some of them.

To facilitate further investigation of changes, there are highlighted links in this table that allow one to: view all the traceroutes for a selected remote host (as a color coded web table accessible by clicking on the node name); access text suitable for attaching to trouble reports (Sum); review the log files (Log*); review the route numbers[I] ("R") seen for a given host together with when last seen; view the available bandwidth time-series for the last 48 hours; to select times and remote hosts for which one wishes to view topology maps, or provide information on the Autonomous Systems (AS) along the routes. In Figure 9, for hours "01" and "14", it can be seen that there were multiple route changes to node1.niit.edu.pk seen from SLAC. Each entry (there can be multiple for each box representing an hour) provides a dot to denote that the route has not changed from the previous measurement. If the route has changed, the new route number is displayed. The first measurement for each day is displayed with its route number. This very compact format enables one to visually identify if several routes changed at similar times, (i.e. route numbers appear in one or two columns for multiple hosts (rows)), and whether the changes occur at multiple times and/or revert back to the original routes.

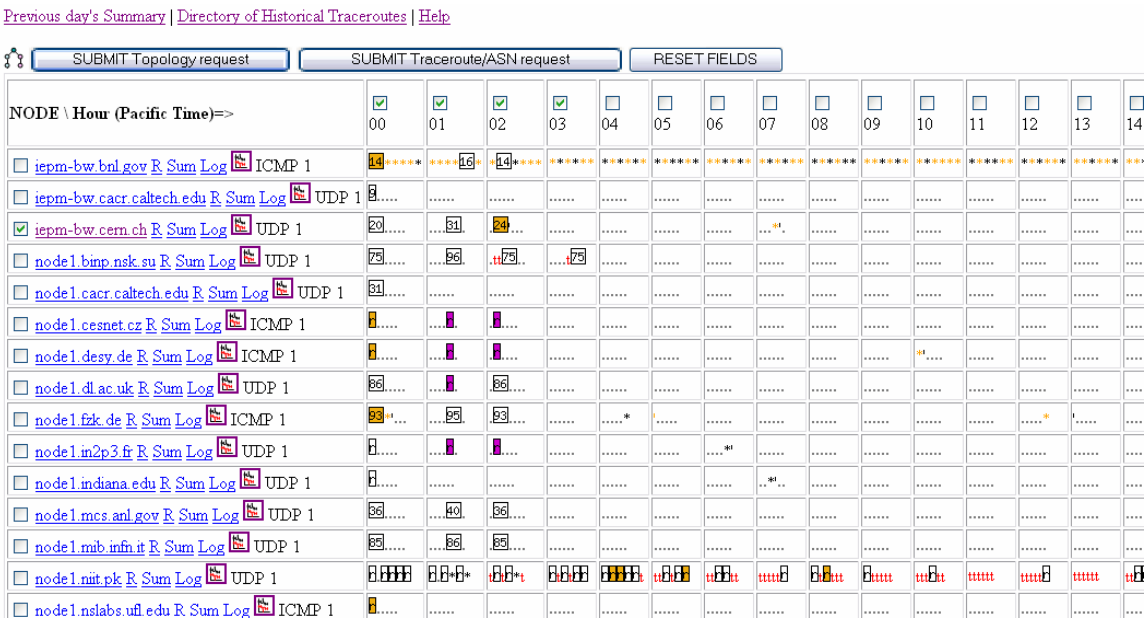

Figure 9 Screen shot of part of a traceroute summary web page with summary table

[I] The route number is simply an index assigned to a unique route between two hosts.

5.2.1 Example of Effect on the Round Trip Times due to Route Change

In this example, the use of traceanal would be illustrated by considering a small route change from SLAC to CERN on August 07, 2005 from 01:56 to 02:16 hrs. Figure 11, is a snapshot of the round trip time (RTT) from SLAC to CERN on August 07, 2005 corresponding to the effect of RTT because of the route change shown in Figure 12 (iepm-bw.cern.ch change from route # 20 to route # 31 and then to route # 24). By selecting the topology request, the graph (Figure 13) of the topology change responsible for the rise in RTT can be obtained.

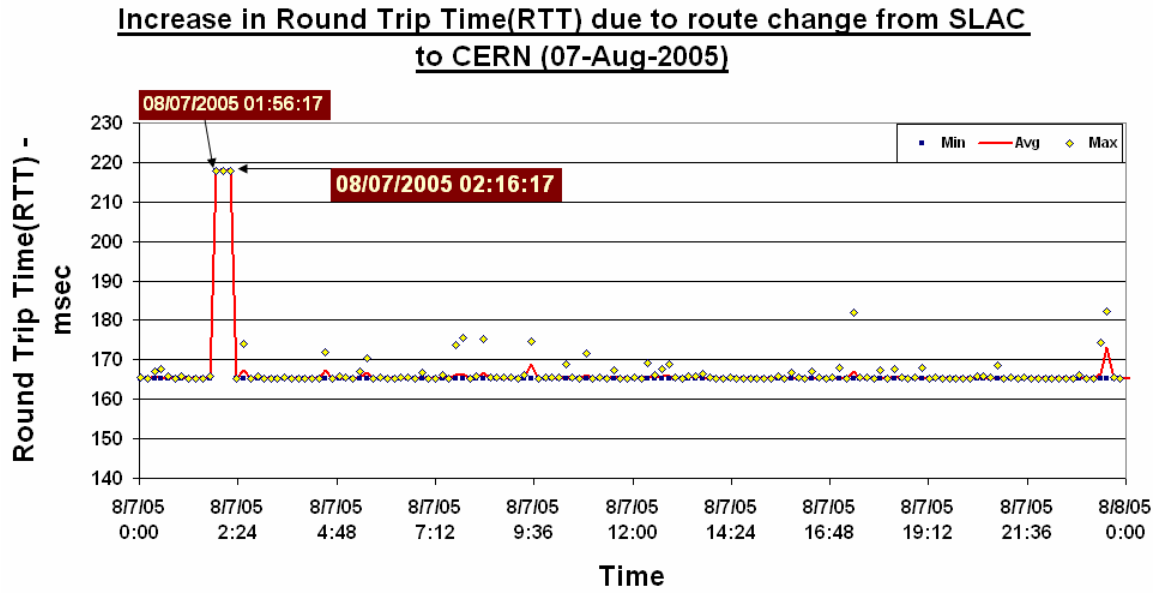

Figure 10 Plot of Round Trip Time from SLAC to CERN

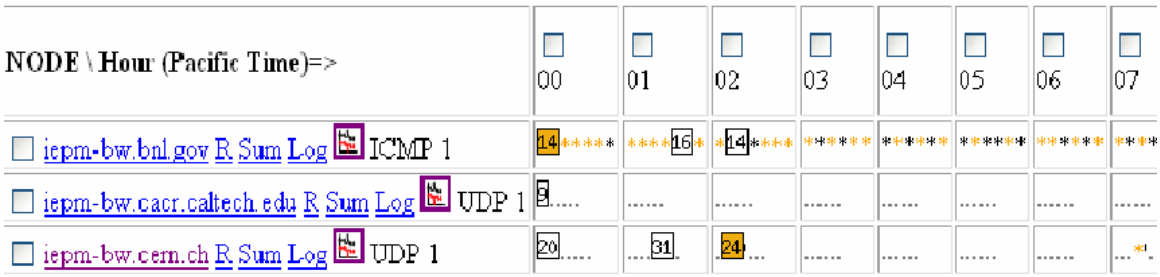

Figure 11 Snapshot of traceroute summary entry for SLAC to CERN on 08/07/05

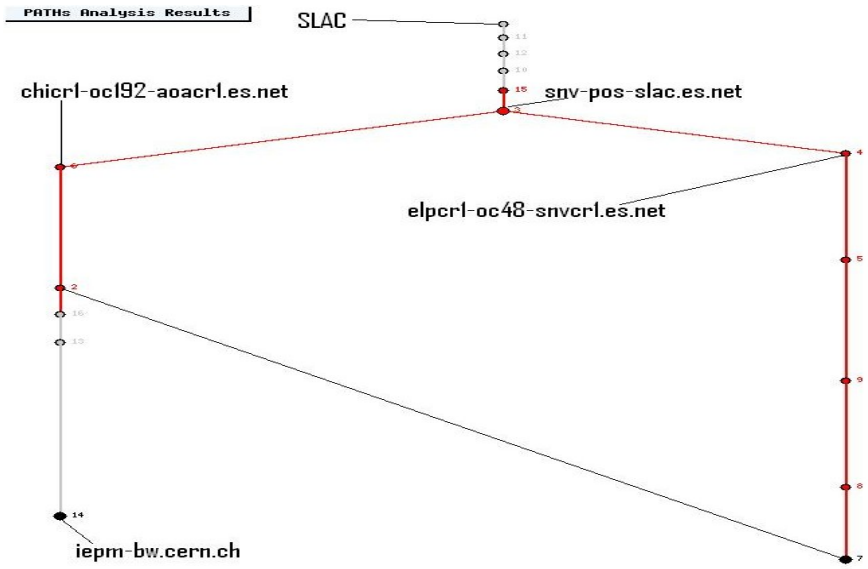

Figure 12 Graphical traceroutes display. Note route change after snv-pos-slac.es.net

5.4 Forecasting and Anomalous Event Detection for IEPM

Network forecasting is a technique to forecast the behavior of the network by "extrapolating" various network measurements e.g. Available Bandwidth, Throughput, Round Trip Time etc. Forecasting enables the system administrators, network analysts and scientists with data intensive needs, to visualize the expected behavior of the network based on the historical behavior of the network. It also enables middleware, such as Grid based data replica selection, to determine the optimal places to get data from, based on forecasts for the near future. Given forecasts, results of the deviations of the forecast, (the expectation) from the observed behavior and anomalies can be obtained.

Anomalies are any significantly abnormal and unusual behavior in the network. They can arise in any network irrespective of their geographical limits and boundaries. Network anomaly detection is the detection of abnormal traffic conditions on a monitored network. Using the potentially overwhelming amount of monitoring data to track down problems is itself a problem. New tools are needed so that users and engineers can quickly spot anomalous behavior or conditions that affect E2E performance. These tools must provide predictive and alerting functionality, with easy drill-down capabilities. Some noteworthy rationales for sudden cropping up of this anomalous behavior in the network are router outages, configuration changes, flash crowd and abuse.

MAGGIE-NS has implemented and compared the use of several different algorithms to detect significant, persistent anomalous events in real E2E Internet performance measurements[24]. The algorithms include the Plateau Algorithm (PA)[25], the Kolmogorov-Smirnov (KS)[26] technique and the use of Principal Components Analysis (PCA)[27] for detecting anomalies. These techniques assume there are negligible seasonal effects in the data (or that the seasonal effects have been removed), To allow for seasonal effects, the Holt-Winters (HW)[28] and Burgess Techniques[29] have also been implemented and compared for making forecasts.

The measurements are based on active probes running on 40 production network paths with bottlenecks varying from 0.5Mbits/s to 1000Mbit/s. To enable comparisons between the various techniques a canonical set of data for 30 sites for 100 days was prepared and studied and anomalous events of interest were identified and categorized. Moreover, the false positives were also identified and missed events for the various methods were categorized. For well behaved data (no missed measurements and no very large outliers) with small seasonal changes most algorithms identify similar events. However, for data with large seasonal changes (e.g. diurnal or weekly), PA or KS can detect false positives especially at the end of a weekend. The algorithms were compared for robustness with respect to false positives and missed events, especially when there were large seasonal effects in the data. The applicability of the algorithms in terms of their intuitiveness, their speed of execution as implemented, and areas of applicability were also studied. Figure 14 shows an example of applying HW to ABwE data measured from SLAC to Indiana. It is seen that the HW forecast tracks the seasonal (diurnal) changes and we can detect the anomalous event at the beginning of Saturday July 31st. This technique is aimed to providing forecasts up to a week in advance and automatically detect anomalies with few misses and few false positives.

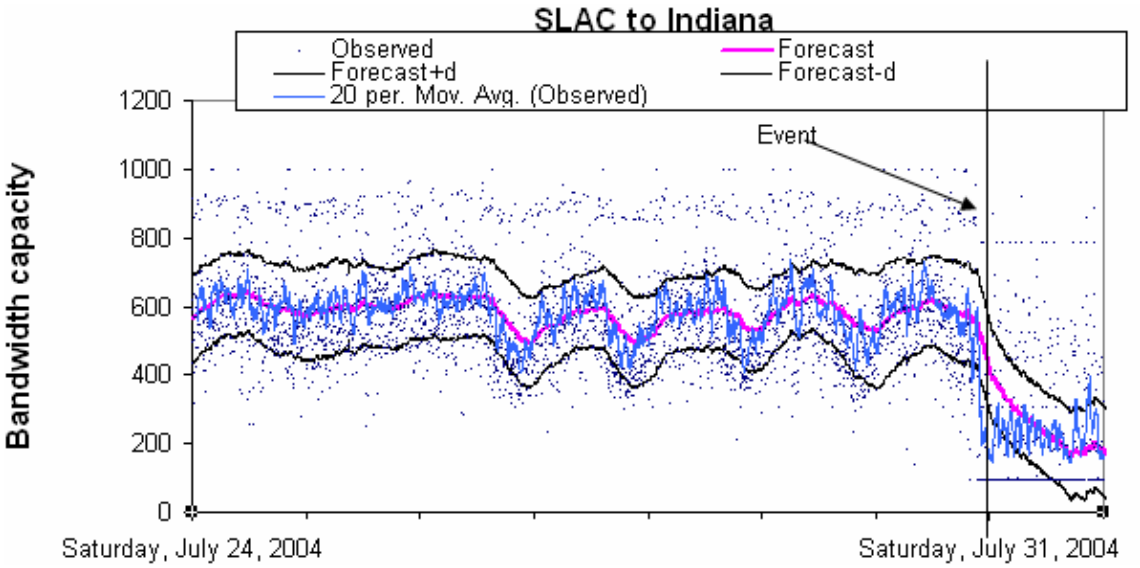

Figure 14: HW forecasts of ABwE data from SLAC to Indiana. The blue dots are the measurements, the blue line shows the smoothed measurements to help guide the eye. The magenta line shows the HW forecast and the black lines show the expected variation.

5.4.1 Active & Passive Measurements

Besides using regular active measurements for forecasting, passive measurements for user data transfers are being investigated. This would add data for more paths, does not add extra network traffic, is related to real applications, needs, and collaborators, and does not require the measurement team to get accounts, passwords or certificates for the remote sites. To make the passive measurements, Netflow[30] records from the border router at SLAC are used. These records are anonymized and filtered to select large (> 100Kbytes), long (> 8 seconds) data transfers from selected applications (ports for iperf, thrulay, bbcp, bbftp, GridFTP, scp etc.) between SLAC and about 30 major collaborator

sites. The flow throughputs are calculated from the recorded start and stop times and the number of bytes transferred.

At first, the passive flow information with that from corresponding active measurements was compared. Typically the agreement was good. The time stamps of the measurements agreed to within 30 seconds and the throughputs to within 3%. In both cases the reasons for the discrepancies are understood.

The distributions of the passive throughputs were reviewed and it was observed that peaks at well-known bottleneck capacity limits (e.g. T1, T3, and Fast Ethernet). Peaks are also observed at corresponding to a combination of RTT and the default TCP window size limiting the throughput. This suggests that there is room for improvement in achievable throughput by optimizing the maximum TCP window sizes.

By aggregating the flows by remote site and folding the data from four week periods onto a single week, 15 remote sites, each with > 1440 flows, i.e. on average a flow each 30 minutes were found. Active monitoring of only 3 of these 15 sites is being done. This would allow adding an extra 12 sites with sufficient data to make forecasts even if there are seasonal effects. About 10% of the sites have significant seasonal changes and for the remaining 90%, less frequent data is needed to make forecasts. Next steps include evaluating the variance of the passive throughput measurements by site and by application.

This project is also investigating the usefulness of Neural Network based approach for forecasting the weather of the network. Initial studies indicate neural network based approaches perform better than or equal to statistical approaches most of the time.

5.5 Using PCA to detect Anomalous Variation in the Traffic

The focus of this research is to detect network anomalies using a technique called Principal Component Analysis (PCA). PCA when tested on data gathered from more than 10 E2E links all over the world proved to be a practical approach for the detection of anomalies and filtering-out patterns of anomalous traffic. This approach was also evaluated in comparison with other techniques such as KS, HW and PA.

Basically PCA is a coordinate transformation method that maps a given set of data points onto new axes. These axes are called the principal axes or principal components. When working with zero-mean data, each principal component has the property that it points in the direction of maximum variance remaining in the data, given the variance already accounted for in the preceding components. By identifying this variance of the data from the average we can detect and perform separation of anomalous data on the link in an efficient and economical manner thus preventing ourselves from false alarms to a very large extent.

5.5.1 Areas of Applicability

PCA is currently being used in many major research domains like bio-informatics, image processing & pattern recognition, data compression, geo-physics and also for the separation of anomalous patterns of data from huge volumes of network data. In short it can be said that PCA is an ideal technique when the behavioral phenomenon of multi-

dimensional data sets are to be analyzed; hence the inspiration to apply and evaluate PCA for the detection of network anomalies in active E2E network measurements.

Using E2E network measurements for the detection of abnormal behavior compared to reading SNMP MIBs from network equipment (see for example reference 27), enables end-users to gather their own data (since reading SNMP MIBs typically requires special privileges). Initial testing shows that PCA not only detects anomalies efficiently but also the analysis resource consumption is favorable compared to the other techniques.

6. Case Studies

6.1 Fiber Outage in Pakistan June 27th 2005 to July 8th 2005

Pakistan's sole under sea optical fiber link, called Southeast Asia, Middle East and Western Europe-3 (SEAMEWE-3), stopped working due to a fault[31] from the 27th June to the 8th of July 2005[32]. This disruption halted the global connectivity of almost 10 million internet users in the country. The PingER team carried out an analysis[33] of this event in terms of its impact on network performance (reachability, loss, RTT etc.) for Pakistan sites as seen from outside Pakistan.

6.2 Pakistan Internal and International Connectivity

Given the Pakistan interest of the MAGGIE-NS collaboration and the need to be able to set expectations for future Pakistan research and education Grid activities, the MAGGIE-NS team also embarked on a more general study of Pakistan's internal and international connectivity. Very interesting results were observed including comparison of the performance of two main ISPs i.e. NTC and Micronet Broadband. This case study provides a lot of information about the congestion in the Pakistan's networks. As a result of this analysis, the NTC staff was provided with solid evidence of their quality of service. The issues raised are now being investigated by NTC. This can be very helpful especially with NTC being the backbone of PERN.

7. Spin-off Benefits

The MAGGIE-NS Project also provided a great opportunity to the research community at SLAC, to identify and understand the issues in the network infrastructure in Pakistan. Regardless of the technical advancements achieved through the project, MAGGIE-NS stands out as a harbinger in raising the standard of research in Pakistan. An institute launched in April 1999 with a vision to *"usher in an Information Culture and Technology Revolution in the Country"*, NIIT has progressed by leaps and bounds, evident by the quick absorption and high standing of the first graduate batch (BIT-I) in the local market. However, it is the "Research Culture" and the growing number of "International Collaborations" that has raised NIIT to one of the best Technology institutes in the country. The NIIT-SLAC Collaboration has certainly played an important role in setting standards for other research groups in the institute.

The NIIT faculty gained many benefits working with the research community at SLAC. Not only did they get a chance to update themselves about the global research in

networks and network monitoring, but the two week visit and constructive critique of the course contents by Dr. Les Cottrell proved to be valuable for the institute. The NIIT faculty got know-how about the latest trends in technology as well the variety of teaching methods followed at Stanford University. This was re-enforced by the visit of Prof. Dr. Arshad Ali to SLAC where he met with senior SLAC and Stanford faculty as well as senior people in successful (Sun and Cisco) and start-up technical industries. Issues discussed included distance learning techniques, training techniques, fostering of technology transfer and entrepreneurship, peace and conflict resolution. Moreover, the Network and System administrators at NIIT were able to see the networking and security policies in place at SLAC, giving them a broader vision of the improvements required in the internal infrastructure at NIIT.

MAGGIE-NS combined the guidance of Dr. Cottrell's broad networking knowledge and the direction of Dr. Ali's vision to transform a group of hardworking and motivated students into dynamic researchers, capable of initiating a revolution, not only in the IT sector in Pakistan, but also other aspects of the society. MAGGIE-NS contributed to the institute by providing the students with an opportunity to spend a year at SLAC and work in a high class research environment, study at a high class university, and interact with high class people, to transform them into high class individuals. In addition, two groups of graduate and undergraduate students located in Pakistan at NIIT have had the benefit of co-supervision by Dr. Cottrell. These experiences will definitely go a long way in impacting the Pakistani society and sparking an ICT revolution in the country.

8. Future

The next proposal to USAID/US-State Department focuses on extending the MAGGIE-NS work to build a comprehensive monitoring infrastructure for the Pakistan Educational Research Network (PERN). In the near future for MAGGIE-NS, we plan to address:

1. Documentation of the new PingER2 installation procedures still needs to be completed.
2. With the growth in number of monitoring sites, policies for determining the beacon sites have to be re-designed, in order to retain the extensive coverage. However, various issues, for instance the impact on the traffic of the hosts in developing countries and probable need for collection of data have to be kept in mind.
3. Several enhancements need to be made to the PingER Management module. More advanced mechanisms for filtering would be developed in future. With the exponential growth in the PingER monitoring data, anomaly detection mechanisms will be developed which identify the anomalies while the data is being analyzed, usually at the end of the day.
4. As the PingER coverage is increasing, various new parameters like Mean-Time-Between-Failure (MTBF) and Mean-Time-To-Repair (MTTR) are being introduced, specifically for the nodes in Pakistan, Africa, India and other

developing and underdeveloped regions. The new parameters would help in identifying the total uptime for nodes in these countries and help in quantifying the reliability of these hosts/paths.

5. Improvements in the PingER visualization are also underway. Modules to generate "Executive Plots" are being developed, specifically for the higher management in the developing regions. This would give them a quick overview of the overall network performance in their regions, as compared to other regions; thus going a long way in decision making.
6. In addition to deploying the landmark on-demand ping measurement tool to the PingER monitoring sites, there are three significant advances in TULIP, which would be carried out after the completion and analysis of the initial version:
    a. The number of landmarks would be increased in order to obtain more accuracy in the actual geographic location of the specified IP.
    b. The value of Alpha would be optimized.
    c. Traceroute analysis would be carried out to study the sudden change in routes, to improve the overall accuracy of TULIP.
    d. The GeoLIM Constraint-Based Geolocation technique will be used to optimize the localization. Once complete, TULIP would have the capability to correctly detect the location of any given IP in any location of the world and effects of network infrastructure would be minimized.
7. The promising study of passive measurements for forecasting still needs considerable work identify how well/consistently it works for various applications.
8. A detailed case study covering various aspects of the Pakistan network performance is being carried out. Classification on the basis of cities, universities, ISPs etc would be made to help identify the most reliable, fast, efficient link the country. Congestion level would be found at various levels i.e. the end host for instance NIIT, ISP, for instance NTC or the core gateway router of the country, for instance the PIE (Pakistan Internet Exchange)[34]. This would help in identifying the exact locations of the bottlenecks, which would in turn help optimize the country's network with much less than projected cost.

9. Conclusion:

The outcome of MAGGIE-NS looks very promising considering the progress and pace of development of PingER and IEPM-BW. The students and researchers at NIIT have contributed a considerable amount of time and taken the lead to help establishing a research culture at NIIT. The initial contribution from NIIT revolved only around trivial tasks, it provided and essential launching pad for this collaborative initiative. The on-going visit of two dynamic individuals to SLAC has certainly increased the pace as well as the nature of these contributions, with significant enhancements in the on-going projects at SLAC. Consequently, the problems in the overall network of Pakistan are being analyzed with a much greater amount of depth with an aim to replicate these

findings to other developing countries. Most importantly, MAGGIE-NS has provided NIIT and Pakistan with a vision that with a little but of help, Pakistan can be self-sufficient in fulfilling its technical needs.

10. Publication & Talks

1. Arshad Ali, Modood Ahmad Khan, S. Azam H. Bukhari and Waqar, ADAM: A Practical Approach for Detecting Network Anomalies Using PCA, SZABIST, ACM and the IEEE Karachi have jointly organized the National Conference on Emerging Technologies (NCET 2004) which took place on December 18, 2004

2. Anomalous Event Detection for Internet End-to-end Performance Measurements, R. Les. Cottrell, Connie Logg and Mahesh Chhaparia, Maxim Grigoriev, Felipe Haro, Fawad Nazir, Mark Sandford. (To Be Submitted)

3. January 2005 Report of the ICFA-SCIC Monitoring Working Group; edited by Les Cottrell for the ICFA SCIC Monitoring Working Group.

4. Forecasting Network Performance presented by Les Cottrell at the Grid Performance Workshop, Edinburgh June 2005.

5. Diagnostic Steps presented by Les Cottrell at the Networking for Non-Networkers second Workshop in Edinburgh, June 2004.

6. Quantifying the Digital Divide from Within and Without presented by Les Cottrell at the International ICFA Workshop on HEP Networking, Grid and Digital Divide Issues for Global e-Science, Daegu, Korea, May 23-27, 2005.

7. Internet Monitoring and Tools presented by Les Cottrell at the International ICFA Workshop on HEP Networking, Grid and Digital Divide Issues for Global e-Science, Daegu, Korea, May 23-27, 2005.

8. Quantifying the Digital Divide from Within and Without; presented by Les Cottrell at the Internet2 Members Meeting SIG on Hard to Reach Network Places, Washington, May 2005.

9. Higher Education and Research in Pakistan presentation to SLAC Computer Services by Prof. Dr. Arshad Ali, NUST, Rawalpindi, Pakistan, April 15, 2005

10. Internet Monitoring lecture given by Les Cottrell at NIIT, Rawalpindi, Pakistan, March 2005.

11. Stanford University, SLAC, NIIT and the Digital Divide dinner talk given by Les Cottrell at NUST convocation, Rawalpindi, Pakistan, March 2005.

12. Datagrid Wide Area Monitoring Infrastructure (DWMI) aka IEPM-BW presented by Connie Logg to the ESCC and the Joint Techs Performance Tools Workshop, February 2005.

13. End-to-end Anomalous Event Detection in Production Networks, presented by Les Cottrell at the ESnet Site Coordinators meeting, Salt Lake City, February 2005.

14. Characterization and Evaluation of TCP and UDP-based Transport on Real Networks, presented by Les Cottrell at the Protocols for Fast Long Distance Networks, Lyon, France Feb 2005.

15. Quantifying the Digital Divide, prepared by Les Cottrell for the Internet2/World Bank meeting, Feb 7 2005.

11. Appendix I; Measurement tools supported by IEPM-BW

11.1. TraceRoute

Traceroute is a very powerful tool for diagnosing network problems. It helps in finding the number of hops to a remote site and how well the route is working. Traceroute was used to identify routing problems from NIIT to SLAC. A very interesting piece of information was revealed about the routing of NTC that sometimes traffic going from Pakistan to Pakistan was going via America.

11.2. Ping

Ping utilities use the Internet Control Message Protocol (ICMP) to determine the availability and responsiveness of network hosts. It sends an echo requests to the address specified and lists the responses received and their round trip time. When the utility is terminated it summarizes the results, giving the average round trip time and the percent packet loss. This utility can be used to determine whether there is a problem with the network connection between two hosts. Ping is the main utility used for the PingER project. This tool also helped in detecting various network problems, including:

- Identifying where hosts that are said to be located in Pakistan are actually in Pakistan
- Studying the overall reliability of connections
- Detecting network anomalies based on Round Trip Time (RTT) and Packet loss
- Calculating bandwidth between links using the above two parameters

11.3. Thrulay

The program thrulay is used to measure the capacity of a network by sending a bulk TCP stream over it. Unlike iperf (see below), thrulay not only reports achievable throughput, but round-trip delay time as well.

11.4. Iperf

Iperf is a tool to measure the bandwidth, allowing the tuning of various parameters. Iperf, as we used it, reports achievable TCP throughput. It is an important component of the IEPM-BW project and has helped a lot in determining network capacities and utilization of these capacities. Due to their ability to congest the network, this tool and thrulay are used less frequently than the others.

11.5. ABwE

ABwE This tool can be used in a relatively continuous mode (e.g. a measurement each second) and detects substantial bandwidth changes caused by improper routing or by congestion. The usefulness of such a tool has been proven several times during last months when we discovered several dramatic routing changes which were corrected soon after we identified and reported them. The NTC's fluttering route problem was detected with ABwE verifying the results of our earlier measurements. It has also been discussed in the case study which has been prepared by the PingER team, and will be discussed towards the end.

11.6. PathChirp

PathChirp is an active probing scheme that uses a novel "packet chirp" strategy to dynamically estimate the available bandwidth along an E2E network path. Internet and

test-bed experiments as well as simulations reveal that pathChirp provides accurate, though somewhat conservative, estimates of the available bandwidth. Though pathChirp uses a factor of ten more time and network traffic to make its measurement, it is more accurate than ABwE. On the other hand it is almost as accurate as pathload[35] and uses a factor of ten less network traffic than pathload. The mix of a low network intrusive packet pair bandwidth measurement tool (such as ABwE or PathChirp) with a more intrusive, user-centric TCP throughput tool (such as iperf or thrulay) appears promising to provide low-impact, short term (updates/minute) real-time measurements that are normalized to less frequent more accurate measurements.

12. References

[1] www.nust.edu.pk

[2] www.qau.edu.pk

[3] *The PingER Project: Active Internet Performance Monitoring for the HENP Community*, Warren Mathews and Les Cottrell, Stanford Linear Accelerator Center, IEEE Communications Magazine, May 2000

[4] http://www.slac.stanford.edu/comp/net/wan-mon/tutorial.html#ping

[5] *Quantifying the Digital Divide from Within and Without* presented by Les Cottrell at the International ICFA Workshop on HEP Networking, Grid and Digital Divide Issues for Global e-Science, Daegu, Korea, May 23-27, 2005

[6] MonALISA: A Distributed Monitoring Service Architecture; H.B. Newman, I.C. Legrand, P.Galvez, R. Voicu, C. Cirstoiu; CHEP 2003, La Jola, California, March 2003

[7] http://www-iepm.slac.stanford.edu/bw/

[8] *pathChirp: Efficient Available Bandwidth Estimation for Network Paths*; Vinay J. Ribeiro et al; ;PAM 2003, San Diego Super Computer Center, University of California, San Diego

[9] *Comparison of Public End-to-end Bandwidth Estimation Tools on High-Speed Links,* A. Shriram, M. Murray, Y. Hyun, N. Brownlee, A. Broido, M. Fomenkov, k. claffy, PAM 2005

[10] *ABwE :A Practical Approach to Available Bandwidth Estimation;* Jiri Navratil and R. Les. Cottrell Stanford Linear Accelerator Center (SLAC); PAM 2003

[11] http://dast.nlanr.net/Projects/Iperf/

[12] *thrulay, network capacity tester*, S. Shalunov, available at http://www.internet2.edu/~shalunov/thrulay

[13] "Super Computing 2005", see http://sc05.supercomputing.org/

[14] www.ntc.net.pk

[15] www.pern.edu.pk

[16] http://www.pern.edu.pk/home/network_diagrams/networkdiagram_files/slide0003.htm

[17] www.hec.gov.pk

[18] http://www.bahria.edu.pk/

[19] *Constraint-based Geolocation of Internet Hosts*; B. Gueye, A Ziviani, M. Crovella, S Fdida, ACM Internet Measurement Conference 2004

[20] *Scale-free behavior of the Internet global performance;* Roberto Percacci_ and Alessandro Vespignani† European Physical Journal B 32 411-414 (2003)

[21] http://www.traceroute.org/#Looking%20Glass

[22] : *Experiences in Traceroute and Available Bandwidth Change Analysis*, presented by Connie Logg, Les Cottrell, and Jiri Navratil at the Sigcomm2004 Net Trouble Shooting Workshop Portland Oregon Sept 3, 2004. Also SLAC-PUB-10518

[23] http://www.slac.stanford.edu/comp/net/iepm-bw.slac.stanford.edu/tracesummaries/today.html

[24] *Cause and Effect of Anomalous Events on QoS Parameters;* Márcio de Freitas Minicz and Alessandro Anzaloni; 2nd International Workshop on Inter-Domain Performance and Simulation (IPS 2004) 22-23 March 2004, Budapest, Hungary
http://w3.tmit.bme.hu/ips2004/papers/ips2004_028.pdf

[25] *Automated Event Detection for Active Measurement Systems*; A. J. McGregor and H-W. Braun, Passive and Active
Measurements 2001. Available at http://byerley.cs.waikato.ac.nz/~tonym/papers/event.pdf

[26] http://www.physics.csbsju.edu/stats/KS-test.html

[27] *"Diagnosing Network-Wide Traffic Anomalies*"; Lakhina A, Crovella M, Diot C; Sigcomm 2004

[28] http://www-iepm.slac.stanford.edu/monitoring/forecast/hw.html

[29] *Principle components and importance ranking of distributed anomalies*; Kyrre Begnum and Mark Burgess; Machine Learning Journal, 58, 217-230, (2005)

[30] http://www.cisco.com/warp/public/732/Tech/nmp/netflow/index.shtml

[31] http://www.jang.com.pk/thenews/jun2005-daily/28-06-2005/main/main3.htm

[32] http://www.jang.com.pk/thenews/jul2005-daily/08-07-2005/main/update.shtml#23

[33] http://www.slac.stanford.edu/grp/scs/net/case/pakjul05/index.htm

[34] www.pie.net.pk

[35] http://www.cc.gatech.edu/fac/Constantinos.Dovrolis/pathload.html