

20180503 PingER Team Meeting

Time & date

Thursday, May 3rd 9 pm **Pacific time**; **Friday, May 4th 2018 9:00 am Pakistan time; 12:00 noon Malaysian & Guangzhou time; and 11 am Thailand time.**

Format

New items and updates are in bold face.

Coordinates of team members:

See: <http://pinger.unimas.my/pinger/contact.php>

Attendees

Invitees:

Wajahat Hussain (SEECs) (may not be able to join); Saqib+; Johari (may not be able to join); Adib?; Fizi Jalil (MYREN); Dr. Charnsak Srisawatsakul+ (Ubru), Les+, Bebo?, Umar-

+ Confirmed attendance

- Responded but Unable to attend:

? Individual emails sent

Actual Attendees

Saqib, Charnsak, Les

Others

Administration

- **Membership of pinger-my** in <https://groups.google.com>.

Amity (Updated 5/3/2018)

- **Their PingER MA is up and running and we are successfully gathering data. We are working on the reliability.**
- **They have made progress on the Android version of the PingER MA.**
 - **They now automatically get the updated Beacon list**
 - **We are looking at how to store the data and set up a proxy**

Bebo (No update 3/8/2018, 5/3/2018)

Looking into moving PingER to a "[blockchain](#)" database good for decentralizing distribution of data. Monitoring sites would then be able to write to a distributed ledger. This would change the architecture to a more peer to peer architecture. It helps with continuity of PingER since reduces dependence on a single site (SLAC). See Block Chain in [Future PingER Projects](#). Bebo sent several references to Saqib who has looked at them. We could start with real-time data without including the whole archive, i.e. in parallel to the continued centrally managed archive. It would be a private Blockchain and hence not be as compute intensive as a public blockchain. Johari is also interested and will follow up with Bebo and Saqib.

Saqib gave an outstanding presentation at Guangzhou. There were a lot of discussions. Saqib has submitted to IEEE transactions and should hear mid May if it has been accepted. The conference is in September. Saqib and Bebo will discuss the next steps. Saqib will try and find a graduate student to work on it.

Thailand (Updated 5/3/2018)

The UBRU IPv4 PingER MA continues to function reliably. Charnsak also has an IPv6 PingER MA working at UBRU. He downloaded the latest version of pinger2.pl and set up a set of IPv6 hosts in <HostList> in pinger.xml. There are a couple of fixes needed to make the node name consistent. At a later stage we will need to figure out how to automatically support IPv6 beacons without messing up MA's that do not have IPv6 capability.

He is using ~ 100 IPv6 targets from Saqib. They are in the<HostList>.

Umar and Charnsak are working on making measurements from UBRU comparing ping/ICMP vs TCP RTTs. Just finished running yesterday after 4 days.

Charnsak is looking at a host in Chan Parsa province in Laos as a potential site for a PingER MA.

Charnsak would like to have write access to parts of the PingER Wiki site. Les investigated and it appears this can be done, even if Charnsak does not have a SLAC account. Les sent Charnsak the relevant information after the meeting.

Les is pursuing the possibility of an MoU with SLAC legal, the legal contact was on maternity leave so there was a delay. He has returned and been reminded and is looking at it. He says "The process for foreign MOUs has unfortunately gotten more complex per new DOE requirements." It can take a long time to go through the process. SLAC has not submitted any MoUs in the last 5 years. Les is still pushing forward and sent an outline to his bosses. **His is not looking hopeful.**

UUM (Updated 5/3/2018)

- Regarding the paper " Socioeconomic Development indices and their Reflection on Internet Performance in the ASEAN Countries ", Adib has submitted the paper to IEEE/ACM Transactions on Networking because IEEE Access does not accept a resubmission of previously rejected paper. Even so, he explained to them that our new submission has a totally new contribution, but they do not take it. The paper status is under review.
- Adib, Bebo, Les met with Southampton Web observatory person. There seemed to be enthusiasm. Adib was going to send some materials to Southampton. The person at Southampton gave us some links. Adib is in the early stages of exploring what web observatory data to link with such as business context indicators, social media and government sites. There is no update 3/29/2018.

NUST: (Updated 5/3/2018))

They have shortlisted candidates for the task of managing the Pakistani MAs. The students will be given a stipend. **As of 4/30/2018, they are still waiting for new students. The interns have not joined the lab yet. So the progress is a little slow.**

Wajahat proposes to get a list of the new Universities in Pakistan and contact them encouraging them to participate in PingER and set up MA. They have made a list of new university sites, communications networks, Labs in different regions of Pakistan (especially the remote regions) and will make contact.

There is an upcoming grant call for projects between Pakistan and the US. Topics may be focused on cybersecurity, health, and education. It has not been announced yet. Wajahat will get the details and share them with the team as soon as they are available. It is interesting since getting a US partner appears to be a roadblock for many potential Pakistani responders. However, the topics may not be very related to PingER. NUST is looking at applying to set up a cyber lab. Getting the funding will be in competition with other Pakistani Universities.

For cyber the main things we could think of from PingER were: quantifying what fraction of hosts block pings, punching holes in firewalls to allow pings, how to misuse ping (e.g. ping-of-death, or using anomalous ping packets to deduce the OS etc. flood pings for DOS), the host can respond to ping but applications do not work. Fear of misuse of pings can result in the system administrator, network administrator or cybersecurity blocking pings. A possibility might be a study of what fraction of say working www/dns etc. apps (i.e. checking if a host responds to the relevant port) do not respond to pings. This could be by application, by country or by region etc. Also how to protect a remote pinger traceroute or server from being used in DOS attacks. As of 3/27/2018 there is no call so far. There was one last year, so Wajahat is expecting one.

We were unable to gather data from:

- 121.52.146.180 (kohat.edu.pk) down since Nov 22/2017. Wajahat recommends continuing at least until the new student is up to speed (3/8/2018). No data available 3/24/2018.
- cae.seecs.edu.pk last time we were able to gather any data was February 27th.
- pinger-ncp.ncp.edu.pk pings but can't gather data 8/11/2017 and 9/16/2017. Contacted. Pings but can't gather data 10/24/2017. They are in the process of restoring 1/17/2018. Still down February 28, 2018, await new student. (3/8/2018). No data 3/24/2018.
- pinger.isra.edu.pk unable to gather data since 3/6/2018, also does not ping.

UNIMAS (No update 3/8/2018, 5/3/2018)

- Johari ran into a problem with the Raspberry Pi image creation. Apparently, the image has to be burnt with exactly the same size as the capacity of the micro SSD, and the latter varies. There may not be a solution. Do we give up?
- Johari is looking at updating the PingER Malaysia website (pinger.unimas.my)
- We have lost both MAs at UNIMAS
 - Johari has been unable to contact Hafiz to get MyREN monitor at UNIMAS (perfsonar-unimas.myren.net.my) working again. There was a discussion between Johari and Adib. Adib confirms Hafiz is still at MYREN, MYREN are moving locations which may have an impact on some servers and availability of Hafiz. Adib will try and contact Hafiz.

UAF/GHZU (Updated 5/5/2018))

He has written a very nice paper on blockchain and its potential use for PingER storage, reviewed by Les and Bebo (also see above under Bebo).

IPv6 node in Beijing is up and running for the last 3 weeks but not available from outside China. Does this host have a name or what is its IPv6 Address (is the latter 2001:da8:270:2018:f816:3eff:fef3:bdf3)? Is it possible to make it accessible from say SLAC (134.79.0.0/16)?

Saqib has run Umar's scripts to compare ping vs TCP. It took about 4 days and Umar has the results.

Saqib has noted that the RTTs from SLAC to China tend to be greater than those from SLAC to S. E. and South Asia. Looking at the monthly avg RTTs from SLAC to: Japan, China, S. E. Asia and South Asia. Les does not see this, it needs further investigation. Below are the RTTs from <http://www.wanmon.slac.stanford.edu/cgi-wrap/pingtable.pl>. The Directivity is described in <http://www.slac.stanford.edu/comp/net/wan-mon/tutorial.html#directness>. It is a metric to identify the directness of the connection between 2 nodes at known locations. Directness values close to one mean the path between the hosts follows a roughly great circle route. Values much smaller than 1 mean the path is very indirect.

Region/Country	25% RTT(ms)	Avg RTT(ms)	Median Rtt (ms)	Pairs	Median Directivity
Japan	110	116	114	6	0.75
China	187	217	201	161	0.52
Malaysia	214	242	236	27	0.58
Thailand	220	229	231	8	0.56
S. E. Asia	217	235	236	83	0.58
Pakistan	282	278	285	14	0.43
South Asia	269	287	286	46	0.44

PingER at SLAC (Updated 5/5/2018)

Umar looking at extending the comparison IPv6 vs IPv4 ping RTTs and TCP vs ICMP/ping RTTs. See [Towards Analysis of ICMP vs TCP Ping Latencies](#).

- See [Towards Analysis of ICMP vs TCP Ping Latencies](#)
 - IPv6 results gathered using [ping-vs-tcp.pl](#) script. They are yet to be analyzed
 - Tests complete in less than a day; not many IPv6 addresses
 - About 56 nodes with IPv6 addresses, 14 of which responded with Npings
 - We essentially have a 14-point data set
 - IPv4 results gathered from SLAC. (Repeating Virginia Tech experiments.)
 - Complete batch may be [downloaded here](#) (approx. 24 MB)
 - Skimmed results; findings are pretty much the same as before
 - Identified relevant events in the network stack that highlight timing (_RECVFROM, _RECVMSG, _IP_RECV, _NETIF_RX etc.). Looking for instrumentation that enables us to measure timestamps. We also need to figure out how to determine whether ICMP & TCP traffic are treated differently? and then how to measure the difference?
 - perf-tools allows us to measure transport events
 - If we could assume that the path for ICMP & TCP through the network is the same, then the only difference between two (controlled) tests would be the time spent in the transport layers. This can be measured using perftools.
 - However, such measurements must be made in a controlled environment where ICMP and TCP are treated the same. (I say so because some results — e.g., in East Asia and South Asia — clearly show that ICMP performs much worse than TCP.)
 - We would also need to cater for cross traffic and queuing delays. Given how small the differences are, one may argue that the variations in measurements are due to cross traffic. Perhaps we should start with controlled tests and then see if real world measurements reflect similar behavior.
 - We need to setup a test environment. We can either setup a bare-metal box or use a VM.
 - I will see if I can arrange for a bare-metal box.
- Want MAs at:
 - SLAC and Virginia tech (Les, Umar) measurements for ping vs nping completed,
 - China (Saqib has agreed to join in),
 - Thailand (Charnsak has agreed to join in),
 - If others wish to join the paper (i.e. make the measurements - takes just over a week elapsed time, help analyze the data and put together the paper), we need to know soon.
 - Pakistan (Wajahat)?
 - Malaysia-Sarawak (Johari)?
 - Malaysia mainland (Adib)?

XSS vulnerability in traceroute.pl

- There is a modified version of the PingER [traceroute.pl](#) CGI script. The new version sanitizes the QUERY_STRING variables to prevent some just exposed cross-site-scripting (XSS, see for example https://en.wikipedia.org/wiki/Cross-site_scripting) vulnerabilities. It is recommended that you replace the existing [traceroute.pl](#) script at your site with the new version, i.e. get (e.g. using wget) [traceroute.pl](#) from <http://www.slac.stanford.edu/comp/net/traceroute/traceroute.pl> and install in your CGI directory (often /var/www/cgi-bin or /usr/local/cgi-bin or /usr/lib/cgi-bin or /var/www/cgi-bin). You may need to make it executable with the command `chmod a+x traceroute.pl`. You can test it by loading the URL <http://yourwebserver/cgi-bin/traceroute.pl>.
- Sent to PINGER-DEV@LISTSERV.SLAC.STANFORD.EDU march 31st, 2018.

PingER IPV6 support

- Working on how to support IPv6 Beacons so it is backwardly compatible, making tests. This is driven by Charnsak's needs for pinger6. cs.ubru.ac.th.
- We will need to add an item to the pingtable.pl form to enable the ability to select IPv4 or IPv6 measurements or both.

Host	State	last seen	Status
pinger-ncp.ncp.edu.pk	pinger-ncp.ncp.edu.pk down	Nov 29, 2018	
121.56.146.180 (pinger.kohat.edu.pk)	Down	Nov 22nd, 2018	
pinger.daffodilvarsity.edu.bd	Does not ping, name resolves, web site not accessible, sent email 4/29/2018, they did something and all is working again	April 21st 2018	Fixed
pinger2.if.ufrj.br	Flaky, down since 4/28/2018, and out of disk space, pinger.xml truncated, email sent 4/28/2017, they are looking at it	April 4/28/2018	
pingeramity.in	It was working, now working on reliability	April 27, 2018	

Next Meeting

Next meeting: Thursday, June 7th 9pm Pacific time; Friday, June 8th, 2018 9:00am Pakistan time; 12:00noon Malaysian & Guangzhou time; and 11am Thailand time.

Old information

GZHU (moved here 3/8/2018)

Saqib submitted a project in CERNET to monitor the performance of IPv6 network using PingERv6. He received the news that the project is accepted with 100K RMB. Now he has 2 accepted projects regarding PingER and total amount he has is near about 40K USD. Further, in his lab, three U1 servers have already arrived through another grant for research purpose. We can also use them for our PingER project.

Therefore, the CERNET has given Saqib a IPv6 based CentOS 6.8 machine in cloud. Now he is trying to deploy the PingER server on the machine. Let's see how it will work on IPv6 based network. This is a 2-year project.

Saqib has made contact with John Pickard author of "Quality of IPv6 Enablement of Universities: An International Study" who has provided a list of about 125 Universities in about 60 countries hosting IPv6 sites. However many are proxies. Les has suggested using perfSONAR (there are about 1000 and they all have lat longs in the perfSONAR database. Saqib is gathering the list, then we will see how many have IPv6 addresses.

The paper title: "Missing Values Imputation in PingER Internet End-to-end Performance Measurements using k-nearest neighbors (k NN)" was not accepted in IMC 2017. He is updating the paper according to the reviewer's comments. Hopefully, Saqib will submit it at some other venue. Not yet decided on the submission venue. Need some suggestions. Updated but not decided where to submit. Update 12/4/2017?

Currently, no data is available on PingER on Android due to unavailability of the live IP address. No update 4/19/2017, 7/6/2017. Email sent to Sara Masood. No update 9/24/2017. Any update 10/24/2017. No progress 1/18/2018.

GZHU (moved here 1/15/2018)

PingER has valuable historical data for the last 20 years. Many analysis and case studies have been carried using this data. A lot of information is available on the website. Saqib's idea is to publish the brief summary all these analysis through a survey paper covering the history and utilization of PingER data starting from 1998 to 2017. Saqib started on it, Les is providing assistance. Need your feedback on the idea of Measuring the Digital Development of the Countries using PingER data. [Is there something you want me to review some, e.g. some draft document on Measuring the Digital Development of the Countries using PingER data](#), or are you asking if it is a good idea to review and create such a document. If the latter I think this is a fascinating subject. Part of the challenge is the chicken and egg problems: i.e. is it network performance influencing advancement of the country, or is it the reverse that advanced countries can afford good networks. My belief is it goes both ways. Also one needs to extend the analysis beyond just Africa else it's kind of a repeat of [Pinging Africa](#), R. Les Cottrell, IEEE Spectrum February 2013. Also see [A Simple Tool for Measuring Digital Development](#), by R. Les Cottrell, IEEE Spectrum February 2013. This is derived from SLAC-PUB-15333.

UUM (moved here 10/24/2017)

"BIND: An Indexing Strategy for Big Data Processing" that uses PingER data. Submitted and accepted by the 2017 IEEE Region 10 Conference (TENCON) that takes place in November. In Penang Malaysia

GZHU

The paper title: "Detecting Anomalies from End-to-end Internet Performance Measurements (PingER) using Cluster Based Local Outlier Factor" is submitted in ISPA 2017 (<http://trust.gzhu.edu.cn/conference/ISPA2017/>). It has been accepted as of 9/17/2017.


The thesis of Aqsa Hameed title "Applying Data Mining and Visualization Techniques on Pinger Data" is published in [ODBMs.org](http://www.odbm.org) and is accessible through <http://www.odbm.org/2017/07/applying-data-mining-and-visualization-techniques-on-pinger-data/>

SEECs (moved here 9/19/2017.)

- Aqsa who was working with Saqib submitted "Applying Big Data Warehousing and Visualization Techniques on pingER Data", Aqsa Hameed, Dr. Saqib Ali, Dr. Les Cottrell and Bebo White, to BDSEA 2016.
- I see it is available from ACM online on the following link: <http://dl.acm.org/citation.cfm?doid=3006299.3006337> for \$15.
- This might be useful to Wajahat's student.

Amity (moved here 9/16/2017)

Preparing a paper on the impact of the cyclone Verdha that hit the Indian coast along with many countries like Thailand, Sri Lanka, Malaysia, Maldives on December 6th. They use K-Means clustering (see https://en.wikipedia.org/wiki/K-means_clustering) to identify anomalies in packets received (inverse of loss) and maximum RTT. Note that for December 7th the reduction in packets received.



Amity (moved here 5/18/2017)

From: Aayush Jain <aayush.2896@gmail.com>

Sent: 24 March 2017 12:31

To: A. Sai Sabitha; harysinha@gmail.com

Subject: PingER Android Team

Abstract for PingER on Android

Progress Made So Far

So far Shivnarayan Rajappa and Rohan Sampson's team have succeeded in making a bare-bones Android Application that can ping beacons, parse data, and generate a text file in a format specified by SLAC ready for uploading. The proposed model involved the application pulling the beacon list from SLAC's servers for pinging. However, the present application has a small percentage of the beacon list hard-coded into the application. As of now, the link between the application and proxy server has not been established.

Future Plans

The new team members are:

1. Rohit Raj
2. Shagun Seth
3. Savy Gupta
4. Aayush Jain
5. Tanuj Saraf

Owing to the advancements in Android technologies during the time of development of the project, our team believes that we can create a more capable and robust application for this project. This involves rewriting the entire application from scratch.

We also propose to create a proxy server that can act as an intermediary between the Android application and SLAC's servers. The proxy server would thus allow handling multiple hosts for greater data collection.

Approach

Our team plans to start off by completing the work on the Android app within 20 days. We will recreate the entire app, with an improved workflow for greater stability. The app will parse the beacon list from SLAC's servers and save as an XML on the device. The data generated after every ping will be appended to a file after cleaning it up with RegEx matches. We first plan to test the app with only a few members of the ping list (which will later be expanded to auto-update in its entirety).

Once we accomplish our work with the app, we will move on to the task of establishing a proxy server. Our entire team will focus on the components of networking, host management, host authentication, file synchronization, and security.

By the end of the project, the server will be able to handle multiple hosts which would all forward it data, and it would in turn reorganise it again for SLAC's servers to pull.

Amity (moved here 4/13/2017)

- The paper on Implementation of PingER on Android has been accepted by IEEE Section. The paper to be online will take 5 months.
- Students are very interested in working with different projects. They have divided the students into three batches (each batch has min of 4 students). The projects currently they are working are:
 - android,
 - data analysis(vardha cyclone)
 - and bigdata

Amity (moved here 3/12/2017)

The students successfully presented the paper on the PingER implementation on [Android.at](#) the confluence 2017 conference.. The paper is submitted to IEEE section.

[Tropical cyclone Vardah](#) hit Chennai in India on the Dec 13th. It impacted the Internet, in particular one of Airtel's undersea cables. Les sent email to A. Sai Sabitha to see if PingER from Amity could see any effect.

- During the next 6 months their research will study the impact Vardha cyclone that hit the Indian coast(South India/Chennai) and a few other neighboring countries in December 2016 as seen by PingER.
 - The idea is to study and analyze the PingER data during the corresponding time frame and deduce significant trends and patterns from the data using
 - 1. Clustering techniques
 - 2. Time series
 - 3. Correlation and Regression concepts

Amity - Java approach (A. Sai Sabitha and Shivanarayan Rajappa)

1. They are using the native java tools, they are not running the [pinger2.pl](#) <<http://pinger2.pl>> script on android since the native java tools have the following advantages
 - a. easier for user,
 - b. no need for prior installation of any software, e.g. load perl interpreter which may require missing skills, especially for a non technical user
 - c. doesn't need a rooted phone
 - d. only the apk needs to be installed to run
2. They have fixed the final sequence number change by using regex, and pushed these changes to github repository.
3. They have installed apache tomcat in the server and plan to use a java file on the server which would connect to the phones that send the request. This java file will then take the input stream received from the phone and write the output stream to a file that would be stored on the server. We are facing some problems regarding a blocked port that is not allowing the phone to connect to the server we are currently working on resolving the issue.
4. SLAC can then regularly pull these files which would be stored based on the month they are received.
5. The Android students have started writing a paper on " implementation of pinger on android " .
6. Next steps:
 - a. Extend the target list by getting the Beacon list from SLAC. It is at <http://www-iepm.slac.stanford.edu/pinger/pinger.xml> on a regular basis and updating the <BeaconList> section at their site. This was part of [pinger2.pl](#).
 - b. Also they will need a utility to clean out old recorded data (say older than 3 months), since it will be gathered from SLAC (via the proxy) and eventually they may run out memory on the Android.

Discussion

To a large extent it depends on how we plan to use this.

- If the phones are just MAs in a fixed location then simply porting [pinger2.pl](#) is easier and probably sufficient.
- If this is intended to grow into a mobile application for general use then it needs to be the Java implementation.

A next step is to get the data from the phone MA to the archive at SLAC. The current method [ping_data.pl](#) requires a public IP address for the phone which may not exist if its is mobile. Getting the MA to put the data to the archive may raise some security issue for the archiver.

Need your feedback on the idea of Measuring the Digital Development of the Countries using PingER data

Two days ago we started being unable to gather data from [pinger.fsktm.um.edu.my](#) (103.18.2.152). When one tries ping it fails,

```
ping pinger.fsktm.um.edu.my
```

```
ping: unknown host pinger.fsktm.um.edu.my
```

```
Exit 2
```

However pinging the IP address works:

```
117cottrell@rhel6-64i:~$ping 103.18.2.152 from http://202.28.194.4/toolkit/gui/reverse\_traceroute.cgi?target=pinger.fsktm.um.edu.my&function=traceroute
```

```
PING 103.18.2.152 (103.18.2.152) 56(84) bytes of data.
```

```
64 bytes from 103.18.2.152: icmp_seq=1 ttl=48 time=265 ms
```

```
64 bytes from 103.18.2.152: icmp_seq=2 ttl=48 time=266 ms
```

```
64 bytes from 103.18.2.152: icmp_seq=3 ttl=48 time=265 ms
```

```
64 bytes from 103.18.2.152: icmp_seq=4 ttl=48 time=265 ms
```

```
^C
```

```
--- 103.18.2.152 ping statistics ---
```

```
4 packets transmitted, 4 received, 0% packet loss, time 3277ms
```

I thought it might be our DNS resolution, however I also cannot see it from Thailand, i.e. from

http://202.28.194.4/toolkit/gui/reverse_traceroute.cgi?target=pinger.fsktm.um.edu.my&function=traceroute

It gives

Can't find IPv4 address for host name pinger.fsktm.um.edu.my. Probably an unknown host.

I get the same result from a host in Pakistan <http://comsatssw.seecs.edu.pk:8080/cgi-bin/traceroute.pl?target=pinger.fsktm.um.edu.my&function=traceroute>