

20181016 PingER Team Meeting

Time & date

Meeting: Tuesday, October 16th 9 pm Pacific time; Wednesday, October 17th, 2018 9:00 am Pakistan time; 12:00 noon Malaysian & Guangzhou time; and 11 am Thailand time.

Format

New items and updates are in boldface.

Coordinates of team members:

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

- Need to add Umar and Dr. Taha - Johari

Attendees

Invitees:

Wajahat Hussain+, Taha+ (SEECs), Saqib+ (GZHU); Johari- (UNIMAS); Adib? (UUM); Dr. Charnsak Srisawatsakul+ (Ubru), Les+, Bebo+, Umar+

- + Confirmed attendance; - Responded but Unable to attend; ? Individual emails sent

Actual Attendees

Wajahat, Saqib, Umar, Les, Bebo, Taha

Others

Administration

- **Membership of pinger-my** is in <https://groups.google.com>.
- Bebo pointed out that the International Web Conference will be in San Francisco in 2019 (see <http://www.wikicfp.com/cfp/servlet/event.showcfp?eventid=77195©ownerid=85335>). There are a whole range of topics. He would love to see something submitted related to PingER.
 - Maybe something on Android PingER and potential uses.
- **Bebo, possible interest in an MA in Jordan - 10/9/2018.**
 - It is a university of Science and Technology.
 - **Bebo is looking to send the Home page:** <http://www-iepm.slac.stanford.edu/pinger/>, plus **Remote Host site invitation letter template**, **Monitor Host site invitation letter template:(txt, doc)**, **form to register a host** to the contact.
 - **Bebo will look into inviting them to our next meeting on Nov 6th.**

Action Items from previous meeting

- **Future meetings are moved to Tuesdays Pacific time.**
- **Figure out how to move forward with the Android PingER project, in particular, schedule a meeting.**
 - **Trying to set a date.**
- **Les: invited Dr. Taha to join the pinger-my email list. Dr Taha joined 9/9/2018. Completed**
- Wajahat: will provide a list of working or potentially working (e.g. being worked on) Pakistani MAs to Les. **They are working on nodes as of email 10/6/2018.**
- **Saqib and Les: worked on being able to gather data from 2001:da8:270:2018:f816:3eff:fef3:bd3 IPv6 node in Beijing working again.**
 - **Les and Saqib added a lot of debugging capability, appears to be successfully working since 9/11/2018. - Completed**
- Umar: update his section on the comparison IPv6 vs IPv4 ping RTTs and TCP vs ICMP/ping RTTs. **Les sent reminder 10/3/2018, see below**
- Johari: enter Umar and Dr. Taha into <http://pinger.unimas.my/pinger/contact.php>, **Les sent reminder email 10/3/2018.**
- Charnsak: We need to get the latest [pinger2.pl](#) measurement agent script installed at Ubru so we can get better logging and see why the other hosts are not being monitored. **Les believes the current logging is sufficient. It would not hurt, in fact, provide a slight improvement, if it is updated. Completed.**

Amity (Updated 8/6/2018, No update 9/6/2018)

- Amity MA is unreliable so using it for a case study does not appear fruitful. **Les is working with Amity to try and understand this unreliability (emails 10/6/2018), they say "There is an internet problem we are taking care of it and will get back soon".**
- The Android version of the PingER MA, is described with comments at [ePingER on Android Native - Amity project](#) (this is a proposal/description from Aayush Jain)

- It describes a multipurpose, stand-alone device that can be widely distributed, something that we have brainstormed about for a long time.
 - We agreed to request Amity to share the App and instructions with us; we will look at installing on a jailbroken Android phone at the San Francisco end and try it out.
 - **They are working on a paper. They sent a 1st draft they have sent to Bebo and Les that we responded to 10/6/2018, a 2nd draft was responded to by Bebo and Les 10/13/2018. It has been submitted to EasyChair.**
 - **They will be attaching the android app for review after testing it again.**
 - **Topher feels that the Android app (when completed and vetted by his team) could easily be installed as a default service on his rainforest monitors (certainly future ones, not devices already in place). Merging the service data that he already collects with that unique to PingER has the potential to lead to some interesting results.**
- Also August 16th and again September 3rd, then **October 6th and October 13th** proposed a meeting between: Bebo, Umar and Les and the Amity folks.
 - 10/24/2018: Proposed: attend a regular one hour meeting of the pingER team: **Next meeting: Tuesday, November 6th 8 pm Pacific time (Nb now on winter time); Wednesday, November 7th, 2018 9:00 am Pakistan time; 12:00 noon Malaysian & Guangzhou time; and 11 am Thailand time, 9:30 am Indian time. The downside is that the focus would be on several PingER matters as well as Amity. However, maybe that is a good way to get started. To do this we would need their Skype ID. I found an ID Sai Sabitha in noida, India, and we appear to have a mutual contact. I left a message.**
 - **The topics are:**
 - Amity's documents describe a multipurpose, stand-alone device that can be widely distributed, something that we have brainstormed about for a long time.
 - Bebo mentioned it to Topher (the Principal Investigator of the RainForest project) and he feels that the Amity PingER/Android app (when completed and vetted by Topher's team) could easily be installed as a default service on his rainforest monitors (certainly future ones, not devices already in place). Merging the service data that he already collects with that unique to PingER has the potential to lead to some interesting results.
 - How should we proceed?
 - We (Topher and the PingER team) agreed to request Amity to share the App and instructions with us; we will look at installing on a jailbroken Android phone at the San Francisco end and try it out.
 - **We also need to finish up the proxy acquisition of data from the Android MAs. A first step is outlined in <https://confluence.slac.stanford.edu/display/IEPM/Proxy+support+for+PingER>.**
 - **A later step is to look at, evaluate and find resources to implement the ideas expounded in the Future Proposals section of the paper you are working on (see below)**
 - Also encourage Amity to put together a paper. **This appears to have been completed.**
 - **Google Firebase Application (see below).**
- The paper has an interesting suggestion for a way forward - **TOPIC to discuss**
 - **Firestore Application:** The need for a proxy server can be completely eliminated by shifting to a cloud-based architecture for managing files. Instead of SLAC pulling in the generated txt files from MAs around the world, the MAs can themselves push these files to a centralized application server hosted on the cloud; which makes it easier for SLAC to access files as per their need. Considering the great degree of integration capabilities Google offers with its products, the flexibility arising from using the Google Cloud Platform would be pronounced. Google's Firestore is a mobile and web application development platform that provides developers with a variety of tools and services to help develop high-quality apps, that can scale easily as per changing demands, and which delivers 99.99% uptime. The Firestore SDK allows mobile app developers to quickly add critical and reliable functionality to their applications in a short time. The recommended option here would be to leverage the Cloud Storage for Firestore that allows robust uploads and downloads onto the Google Cloud Storage buckets. Apart from the user authentication module that comes bundled with Firestore, developers can also declare file security parameters so as to allow only certain file types to be uploaded. Having all user uploaded files in one place will then enable SLAC to access and process files as and when needed. Server-side processing can be done on the Google Cloud Platform as well, thereby eliminating any need for SLAC to maintain its own physical infrastructure. This form of implementation can be highly beneficial in any kind of region of the world, also including remote places, like deep inside tropical rainforests, or places that have recently been hit by a natural calamity. As this implementation model is not dependent on any local measuring agent, the android mobile apps can directly deliver data to SLAC over any form of internet connection in minimal time with high reliability.
 - Bebo agrees it is another reduction on independence from SLAC.
 - It is unclear whether it would work for China since Google services are not available there.
 - Is it only free for low utilization.
 - **Might work for a cache for the latest data, but not for all archived PingER data (hundreds of GB).**
 - It was developed for Android applications, might be good for Android/PingER
 - There may be traffic limits.
 - Umar suggests try for one node and see how it works out in practice

Thailand (No update 8/9/2018, No update 9/6/2018)

Looking at http://pinger6.cs.ubru.ac.th/cgi-bin/ping_data.pl?admin_func=pinger_xml the IPv6 monitoring site pinger6.ubru.cs.ac.th has a pinger.xml configuration file with over 160 IPv6 targets. An example is:

```
<Host>
<IP>2001:630:22:1004:800:ff:fe00:27</IP>
<Name>perfsonar-bw.tier2.hep.manchester.ac.uk</Name>
</Host>
```

Also looking at the data extracted from http://pinger6.cs.ubru.ac.th/cgi-bin/ping_data.pl?in_form=1&begin_hour=00&begin_min=00&begin_sec=00&begin_day=06&begin_month=10&begin_year=2018&begin_offset=&begin_point=y&end_hour=23&end_min=59&end_sec=00&end_day=07&end_month=10&end_year=2018&end_offset=&end_point=y we see ~100 different IPv6 addresses. The example is:

```
pinger6.cs.ubru.ac.th 2001:3c8:d109:1300:202:29:20:125 perfsonar-bw.tier2.hep.manchester.ac.uk 2001:630:22:1004:800:ff:fe00:27 100 1538785462 10 10 199.341 199.525 199.980 1 2 3 4 5 6 7 8 9 10 199 199 199 199 199 199 199 199 199 199 199 199
```

However, according to http://www.wanmon.slac.stanford.edu/cgi-wrap/pingtable.pl?file=average_rtt&by=node&size=100&tick=monthly&from=TH.UBRU.CS.AC.PINGER6&to=WORLD&ex=none&only=all&ipv=6&dataset=hep&percentage=any (which uses the information in the Oracle NODEDETAILS database) there are only about 13 IPv6 targets responding. For example perfsonar-bw.tier2.hep.manchester.ac.uk is missing

It transpires that pinger6.cs.ubru.ac.th has many IPv6 nodes that are not in NODEDETAILS.

We looked in more detail at the first 53 hosts in the <HostList> tag in pinger6.cs.ubru.ac.th's pinger.xml files (see <http://www.iepm.slac.stanford.edu/pinger/pingerworld/pinger6-ubru.txt>). Some of the nodes are already in NODEDETAILS. I have added a few hosts with IPv6 addresses to NODEDETAILS.

Many of the nodes have a single name and both an IPv4 and IPv6 address. Since the key of each record in NODEDETAILS is the name and a name can currently only have one IP address, where we already have the name in the database (with an IPv4 address) I have currently left the record as is.

TOPIC to discuss:

We need to think about how to handle these hosts with both IPv4 and IPv6 addresses. We could do nothing, the simplest solution. Or we could move towards moving to the future and simply change the IPv4 addresses to IPv6 addresses. In a few cases, we may be able to find two hosts at a site or two names for the same host (i.e. an A and a CNAME record), using an IPv4 address on one and an IPv6 on the other. This is the next simplest solution. Doing this loses having a historical record of a target with both an IPv4 and IPv6 address. Umar and Les have some scripts that enable comparisons to be made between access to a host via its IPv4 and its IPv6 address without using the PingER data. An alternative would be to also have a pseudo name for such ambiguous hosts, another would be to modify the database schema. Both of the latter two would require changes to the code in several places. Another alternative might be to create a second record using either the IPv4 or IPv6 address as the name. This requires no changes to the code but is pretty unpleasant. If one could find a CNAME as well as the regular name (A reg) then one could use the CNAME for one of the addresses and the regular name for the other address.

- Looking at a [spreadsheet of 155 perfSONAR hosts](#) from Saqib, Les found 3 that have a CNAME. From these Les has added perfSONAR nodes [planchet.heanet.ie](#) and [perfsonar.ndsu.NoDak.edu](#) to NODEDETAILS the PingER meta database using the A record with the IP address and the CNAME with the IPv6 address.

Charnsak is looking at a host in Champasak University, Chan Parsa province in Laos as a potential site for a PingER MA. Charnsak just got approved to make contact with the Champasak University. He expects to set up the MA in the next 4-5 months (say towards end 2018). It also depends on the partner university, and there may be a lot of paperwork.

UNIMAS (No update 8/9/2018, no update 9/6/2018, no update 10/7/2018)

Need to add Umar Kalim to <http://pinger.unimas.my/pinger/contact.php>. From the 7/5/2018 meeting: Johari can't ssh into the server so he will go to it on Monday. He will also upload the new UNIMAS PingER website next week.

Sent reminder emails 8/6/2018. 9/3/2018.

UUM (No update 8/9/2018, No update 9/6/2018)

Les has sent Adib updates to Figs 3, 4, 5 to extend out to 2018. This is for the paper Socio-economic Development Indices and Their Reflection on Internet Performance in ASEAN Countries

Adib will submit the paper to World Development: <https://www.journals.elsevier.com/world-development>

NUST:

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.

- The list of new universities is ready. Just need resources to engage them.
- Les pointed out that several Pakistani MAs have been disabled, i.e. the data is no longer gathered from them.

- Wajahat will provide a list of working or potentially working Pakistani MAs to Les, and if they are not enabled, Les will enter or re-enable the gathering of data from these MAs.
- **From Wajahat 10/6/2018. "We are working on the nodes you have mentioned. The intern is new and trying hard to understand the programs. Hopefully within a week he will start repairing the nodes."**
-

SLAC was 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. **Disabled 10/3/2018. Wajahat emailed 10/12/2018.**
 - pinger.kohat.edu.pk status is that the concerned person is on leave and the other person doesn't know any status of the system at site.
- cae.seecs.edu.pk last time we were able to gather any data was February 27th. **Disabled 8/19/2018. Wajahat emailed 10/12/2018.**
 - cae.seecs.edu.pk status is that the system at node is totally out of order. The system isn't in working condition. The concerned person there requires full technical support on how the linux can be installed and used. He knows windows only. He also requires all the support on making the system run the pinger scripts etc.
- For several months we were unable to gather data from pinger.isra.edu.pk starting 3/6/2018, also it did not ping. **Then it started working as of September 2nd, 2018 and we gathered data for 3 days. Failed again starting on Sept 5th. At this time (Sept 6th) it was pinging, however, all pings failed (i.e. no response after 31 pings). Now (10/14/2018) it is not pingable.**
- **Wajahat now has a permanent staff engineer so that person can be trained on how to manage MAs and possibly improve the installation process**

UAF/GHZU (Updated 8/9/2018)

Saqib's future at GZHU **will be much clearer after November**, current contract expires February 2019.

Gathering data is working since 9/11/2018.

Blockchain

- 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 BlockChain 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.
- There was a meeting to discuss blockchain possibilities, see [20180709 PingER Meeting on Blockchains](#).
- Bebo's impression is that Saqib will lead in putting the ideas in his paper into practice. Saqib will need some students. Saqib is OK with this. He has 2 masters students but they are working in different areas. Maybe NUST can assist with this. Saqib's partner gave a talk/paper on work so far at the New York meeting on July 31st. The talk went fine but there were not many comments/questions.
- Saqib is pursuing PingER and Blockchain. He is looking at different references shared by Prof. Bebo and the implementation details using Hyperledger Fabric. Saqib is looking at making a test implementation. The blocksize will be 2MB-10MB. It does not appear to be computationally expensive. He will start testing with Internet of Things measurements such as humidity and temperature
- There was a discussion on the use of Databases and whether they could be avoided by caches. Hyperlogic keys are not in an SQL DB, basically, it appears like a cache. There was a question whether a 10MB block would be adequate for PingER. For example, PingER from SLAC has about 700+ targets, the measurements are each 30 mins (48/day) and for pings of 100B and 1000 Bytes i.e. 2*48*700 measurements and each measurement is ~ 140Bytes, so a day's sets of measurements from just SLAC is ~ 10MBytes. We could choose to ignore the 1000Byte pings which would reduce it to ~ 5MBytes/day. The latency of retrieving a block is proportional to the block size. Things will be clearer after the test set up is in use.

IPv6 measurements

- There are now several months of IPv6 PingER measurements from GZHU/BJ, UBRU and SLAC. It is time to think about in-depth analysis of the data.
- Umar has not made much progress on the comparisons of IPv6 vs IPv4 or TCP vs ICMP.
- The PingER measurements would provide longer time spans where one might look for changes with time such as diurnal changes, impact of holidays, anomalies etc.

PingER at SLAC

PingER IPV6 support

- Les identified and sanitized a Cross Site Scripting (XSS) exposure in ping-data.pl. **He sent out a notice to update. Saqib has updated. Unfortunately cannot see from outside the Great Wall of China.**
- Les cleaned up the Beacons. **Now have 1 or more Beacons in > 158 different countries, all are working except one in Syria.**
- There was a problem gathering data from [2001:da8:270:2018:f816:3eff:fe3:bd3](#) the Guangzhou IPv6 host in Beijing. Since it does not use the standard way of gathering data (ping_data.pl), but instead uses anonymous FTP, the standard way of debugging the problem does not work. **Les and Saqib worked on to provide SLAC access to various log files etc. via the anonymous FTP proxy. Added extra debugging to proxy gatherer, it is in production.**

Raspberry Pi version 1 PingER MA at SLAC

- This was installed in mid 2015 as part of the ePingER project (see [ePinger Project at SLAC](#)). It has run flawlessly without manual intervention for over 3 years, automatically recovering from power outages etc. The OS failed a security scan. Rather than upgrade etc., since it has served its time, we decided to decommission it and focus more on the PingER/Android project. Bebo I have the Raspberry Pi and SD card, if you want it.

Host	State	last seen	Status
pinger.cern.ch	Pingable, web server does not respond. Sent email. Bad disk. Working again 9/17/2018.	9/11 /2018.	Fixed 9/17/2018
pinger.cs.ubru.ac.th	Access to http://pinger.cs.ubru.ac.th/cgi-bin/ping_data.pl from SLAC forbidden starting October 1, 2018. Charnsak found that the raspberry had been moved. He restored it to working order 10/4/2018.	9/30 /2018	Fixed 10/4/2018.
nsi.ampath.net	Unable to gather data. Responds to ping and http://nsi.ampath.net/cgi-bin/ping_data.pl , however, no new data is available. Site responded and it is working again as of 10/3/2018.	9/24 /2018	9/30/2018, reminder 10 /3/2018. Fixed 10/3 /2018.
pinger.stanford.edu	Unable to ping the host. Due to building power maintenance.	9/21 /2018	10/3/2018 Fixed
pinger.daffodilvarsity.edu.bd	Unable to ping host, email sent 10/14/2018	10/8 /2018	10/14/2018 fixed
pingersonar-um.myren.net.my	No response	6/26 /2018	
424.62.146.180 (pinger.kohat.edu.pk)	Down-Disabled 10/3/2018	Nov-22nd-2017	
pinger.isra.edu.pk	Down, it came up Sep 2-4, 2018. It is pinging, however, all the targets are not responding after 31 tries, email sent to Wajahat 9/18/2018. As of 10/14/2018 it is not pingable.	March 6, 2018	
pingeramity.in	It has been working since 28th July. It is unclear how stable it is. It is down again 9/5/2018. Worked again from 9/24/2018 to 9/28/2018. Emailed Amity requesting clarification 10/3/2018. They are looking at it.	April 27, 2018	

Context:

- Is there any statistical difference between ICMP and TCP Ping? The context here is the Internet (not data center). This is important because the network stack is different (e.g., MPI over infiniband) and latencies are significantly less.

Questions:

- Why should we focus on minimum RTT instead of average RTT
 - Min RTT essentially reflects fixed delay, while average RTT subsumes variations and path load
- Are the R plots generated using minRTT?
 - Averages and computed. Min RTT is available. Scripts need to be updated to use minRTT.
- What is the breakdown of latency between endpoints? If there is a difference, is it because of the type or location of the source? What if the source of traffic was not SLAC?
 - Latency for an echo packet to travel up the stack and back down is about 3.75 micro seconds (see StackMap <https://www.usenix.org/conference/atc16/technical-sessions/presentation/yasukata>). As expected, this is negligible when considered with milli second latencies. The remaining components would be propagation and queuing delay. As we cannot breakdown the two in a public network without using an active look like
 - To replicate use system tap. See: <https://unix.stackexchange.com/questions/419449/how-can-i-determine-if-a-latency-is-due-to-a-driver-or-the-scheduler>
 - pathchirp etc, we'll continue to consider these as a single component.
- Is there a correlation with the distance between the endpoints?
- Are the differences limited to a particular region? How do we determine/understand if traffic prioritization is implemented?
 - It may be that end hosts which are farther away have larger variances and thus the pronounced differences.
- Test in a controlled environment to avoid variables such as traffic prioritization, queuing delay due to cross traffic.
- Review the time series of latencies for both ICMP and TCP ping, instead of averages?
- Is there a difference between IPv4 measurements vs IPv6.

Next Meeting

Next meeting: Tuesday, November 6th 8 pm Pacific time (Nb now on winter time); Wednesday, November 7th, 2018 9:00 am Pakistan time; 9:30 am India time; 12:00 noon Malaysian & Guangzhou time; and 11 am Thailand time.

Old information

Umar looking at extending the comparison IPv6 vs IPv4 ping RTTs and TCP vs ICMP/ping RTTs. Last update 6/7/2018

- See [Towards Analysis of ICMP vs TCP Ping Latencies](#) - Umar
 - Looked into Traffic Differentiation - Rate Limiting vs. Traffic Prioritization (QoS)
 - Conclusions:
 - Min RTT essentially reflects fixed delay, while average RTT subsumes variations and path load
 - TCP Segment drops manifest as large increases in delay
 - QoS can be implemented in at least two ways:
 - When priorities are implemented, ICMP packets will only be dropped if ICMP quota is full and link is congested, otherwise ICMP traffic is allowed to go beyond quota

- When rate-limits are implemented, ICMP packets will be dropped when ICMP quota is met, even if links are not congested
 - Path loads dictate latency estimates: With weighted QoS implementations, loaded links depress ICMP
- Found relevant papers and technical reports. The intent here is to understand the components of latency to help with the evaluation.
 - Y. Zhang, N. Duffield, V. Paxson, S. Shenker, [On the Constancy of Internet Path Properties](#), in ACM SIGCOMM Internet Measurement Workshop 2001
 - A. Acharya, J. Saltz, [A Study of Internet Round-Trip Delay](#), Univ of Maryland, Tech Rep. CS-TR-3736
 - M. Allman, V. Paxson, [On Estimating End-to-End Network Path Properties](#), in ACM SIGCOMM 1999
 - V. Paxson, G. Almes, J. Mathis, [Framework for IP Performance Metrics](#)
- Setup NeuBot and experimented with tests but wasn't able to find useful examples of traffic differentiation (see <https://www.measurementlab.net/tests/>)
- Skimmed Glasnost to understand traffic shaping (see <http://broadband.mpi-sws.org/transparency/glasnost.php> and <https://github.com/marcelcode/glasnost>)
- [Not Relevant] Miscellaneous notes on ICMP Traceroutes, MPLS tunneling & ICMP (see <http://cluepon.net/ras/traceroute.pdf>), [Measuring Performance](#)
- IPv6 results gathered using [ping-vs-tcp.pl](#) script. About 56 nodes with IPv6 addresses, 14 of which responded with Npings
- IPv4 results gathered from SLAC and Virginia Tech
 - SLAC's batch may be [downloaded here](#) (approx. 24 MB)
 - Skimmed results; findings are pretty much the same as before
- Pending
 - Identified relevant events in the network stack that highlight timing (_RECVFROM, _RCVMSG, _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.

Discussion item (7/5/2018)

Saqib sent an email to the team:

In early years of PingER, the framework was designed to check the latency and other Internet performance metrics between CERN and SLAC to facilitate the data transfer between the two sites.

"I am thinking, is there any possibility to use PingER to monitor the health of the Bitcoin blockchain network? Since the latency is critical in Bitcoin blockchain network as all the incentives depend on the propagation of transactions and mined blocks. Thus, I am only interested in measuring the latency to check its effect on the propagation of the transactions and mined block among different mining pools. I think if we can do such thing on a historical basis as PingER already does for the Internet, it will increase the worth of the framework and its usability.

Maybe a few test experiments can guide us to a good research paper. I am not sure about the feasibility idea, therefore, need your kind feedback"

Umar responded:

"This is an interesting idea. I would like to think about it a bit more though before I respond at length.

From what I gather, nodes in the network may appear and disappear without notice. Therefore, it is safe to assume that the network state is changing at all times. I wonder, what is it that we would be measuring, for it to be meaningful. Would it be the latency between full nodes? Would it be the latency from a PingER monitoring node to the full nodes? Would the monitoring nodes be representative of typical clients? Is latency the metric to measure? What are all the projects that measure latency or other metrics? (The bitcoin nodes project is interesting, which shows the size of the network. Similarly, the project that measures average-transaction-completion time is relevant. As you pointed out, the bitcoin stats about data exchange are relevant too.)

Thank you for sharing the URLs. I will think about it a bit more and get back to you."

Bebo emailed:

- *More and more I am convinced that we could identify an interesting and relevant project taking advantage of our PingER experience and involving Blockchain technology. I also think that it may be possible for us to get funding/grant money for such a project which might clearly draw the attention of Blockchain entrepreneurs, conferences, and publications. It might also allow us to attract the collaboration of other institutions and/or universities that might otherwise have not been interested in the basic PingER goals and technology.*

Though the transactional latency is important especially for achieving consensus, most of the latency is computational that is much greater than the communications latency. Possibly PingER RTT could be included in a measure of BlockChain health and status. Bebo sent a reference (<https://mastanbtc.github.io/blockchainnotes/consensustypes/>) pointing out milliseconds may make a difference to miners to get rewards which is important for BlockChain cash

Saqib sent email of an analysis of the impact of various RTTs, see <https://ethereum.stackexchange.com/questions/18201/does-network-latency-significantly-affect-mining-rewards>.

A possible problem is finding the IP addresses of the bitcoin miners. Bebo dug up a URL showing how it can be done, it is at <https://bitcoin.stackexchange.com/questions/66260/finding-ip-addresses-of-all-bitcoin-miners>.

Bebo's impression is that Saqib will lead in putting the ideas in his paper into practice. Saqib will need some students. Saqib's boss is going to the NY City meeting.

GZHU China - Saqib (moved here 7/2/2018)

Saqib submitted the Camera ready paper on "A Blockchain-based Decentralized Data Storage and Access Framework for PingER" and it has been accepted in Trustcom2018.

Two previously accepted papers are now online on the following links.

- The paper titled "Detecting Anomalies from End-to-end Internet Performance Measurements (PingER) using Cluster-Based Local Outlier Factor" is online on IEEE portal. <https://ieeexplore.ieee.org/document/8367380/> . This paper is not available in SLAC repository. Les will follow up.
- The paper titled "Internet Performance Analysis of South Asian Countries using End-to-End Internet Performance Measurements" is online on IEEE portal. <https://ieeexplore.ieee.org/document/8367431/> . This paper is available in SLAC repository (<http://www.slac.stanford.edu/pubs/slacpubs/17000/slac-pub-17205.pdf>). However, without doi no and other necessary details. Les will follow up.

NUST (moved here 6/29/2018)

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. Emailed Wajahat 6/3/2018 asking for update. He responded "There is no call yet. There was one last year. May be change in the US-Pak policy. Just a guess."

UUM (moved here 6/29/2018)

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 was no update 3/29/2018, or 5/3/2018.

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.odbms.org/2017/07/applying-data-mining-and-visualization-techniques-on-pinger-data/) and is accessible through <http://www.odbms.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 Shivrinarayan 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 it is mobile. Getting the MA to put the data to the archive may raise some security issue for the archiver.

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://comsatsswl.seecs.edu.pk:8080/cgi-bin/traceroute.pl?target=pinger.fsktm.um.edu.my&function=traceroute>