

ePinger Project at SLAC

Table of Contents

- [Table of Contents](#)
- [Introduction](#)
- [PingER](#)
- [Raspberry Pi Model and Specifications](#)
- [Operating System](#)
- [Installation of PingER2 MA on Raspberry Pi](#)
- [Requirements](#)
- [Methodology](#)
- [Measurements](#)
 - [Example target = pinger.unimas.my \(~220 msec.\)](#)
 - [Example target sitka.triumf.ca \(~22msec.\)](#)
 - [pinger.slac.stanford.edu and pinger-raspberry.slac.stanford.edu](#)
 - [Using full set of pings for RTT frequency distributions](#)
 - [Frequency distribution for absolute interpacket delays.](#)
 - [To sitka.triumf.ca from SLAC](#)
 - [From sitka.triumf.ca to SLAC](#)
 - [To ping.cern.ch from SLAC](#)
- [Results](#)
 - [IPDV](#)
 - [Localhost comparisons](#)
 - [Kolmogorov-Smirnov Test](#)

Introduction

This document is the basis of a [paper \(SLAC-PUB-16343\) comparing PingER performance on a 1 U bare metal host to that on a Raspberry Pi](#). It has been running successfully since July 2015.

This is a project suggested by Bebo White to build and validate a [PingER](#) Measurement Agent (MA) based on an inexpensive piece of hardware called a Raspberry Pi (see more about [Raspberry Pi](#)) using a linux distribution as the Operating System (see more about [Raspbian](#)). If successful one could consider using these in production: reducing the costs, power drain (they draw about **2W of 5V DC power** compared to typically over **100W for a deskside computer** or 20W for a laptop) and space (credit card size). This is the same type of power required for a smartphone so appropriate off the shelf products including a battery and solar cells are becoming readily available. Thus the Raspberry could be very valuable for sites in developing countries where cost, power utilization and to a lesser extent space may be crucial.

PingER

PingER (Ping End-to-end Reporting) is the name given to the Internet End-to-end Performance Measurement (IEPM) project to monitor end-to-end performance of Internet links. It is led by [SLAC](#) and development has included [NUST/SEECs](#) (formerly NIIT), [FNAL](#), and [ICTP/Trieste](#), together with UM, [UNIMAS](#) and UTM in Malaysia. Originally, in 1995 it was for the High Energy Physics community, however, this century it has been more focused on measuring the Digital Divide from an Internet Performance viewpoint. The project now involves measurements to [over 700 sites](#) in over 160 countries, and we are actively seeking new sites to monitor and MA sites for this project, as well as people interested in our data. It uses the ubiquitous [ping facility](#) so no special software has to be installed on the targets.

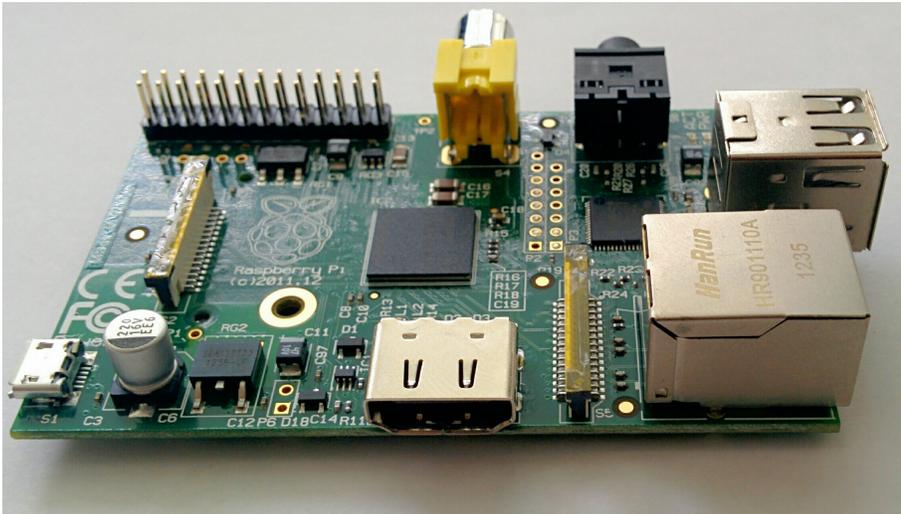
Measurements are made by ~60 measurement Agents (MAs) in 23 countries. They make measurements to over 700 targets in ~ 160 countries containing more than 99% of the world's connected population. The measurement cycle is scheduled at roughly 30 minute intervals. At each measurement cycle, each MA issues a set of pings to each target, stopping when it receives 10 ping responses or it has issued 30 ping requests. From each set of pings one can derive various metrics such as minimum (Min) ping Round Trip Time (RTT) response, average (Avg) RTT, maximum (Max) RTT, standard deviation (Stdev) of RTTs, 25% probability of RTT, 75% probability of RTT, Inter Quartile Range (IQR,), loss, and reachability (host is unreachable if get 100 % loss).

The data is publicly available and since the online data goes back to January 1998, it provides 19 years of historical data on Internet performance.

Raspberry Pi Model and Specifications

The Raspberry belongs to Bebo White and it is the [version 1 of Raspberry Pi, model B](#). The cost is about \$25/each + costs of the SD card. The Raspberry purchased has 512MB RAM, on a 700Mhz ARM CPU and a [32GB SDHC Card](#) (purchased in June 2015 for \$18) was used. It has 2 USB ports, 1 100Mb/s Ethernet interface and 1 HDMI port. For reasons of economy it does not have a Real Time Clock (RTC). Instead, the Pi is intended to be connected to the Internet via Ethernet or WiFi, updating the time automatically from the global [ntp](#) (network time protocol) servers (see <https://learn.adafruit.com/adding-a-real-time-clock-to-raspberry-pi/overview>).

The Voltage requirement for the power is 5V+-5%. Keep in mind that it is necessary to have a keyboard, a mouse and a HDMI monitor to do the installation process, but once PingER is working they are not necessary anymore. We measured the power (Wattage) during normal use and it is 2.7 Watts. When using the Dell mouse, with an LED, powered from the Raspberry Pi it crept up to 3.2Watts. If we have a 5000mA-hour fully charged solar cell battery, then the Raspberry Pi might be expected to run for ~ 25 hours. However we need to be able to recharge from the sun. If we had a [10W 5V solar panel](#) then with an [efficiency factor of 2.5 we may estimate the recharge time](#) of the battery to be ~6.25 hours.



Operating System

The installed system is called [Raspbian](#) a Debian Linux variant. The OS had Perl, Make, dig, ping and mail installed. We accessed it through the graphic interface of [Raspbian](#). We just had to install: [Apache](#) and [XML::Simple](#).

Installation of PingER2 MA on Raspberry Pi

The first step, before start the installation process we had to [change the hostname in Raspbian](#).

```
sudo nano /etc/hostname  
  
sudo /etc/init.d/hostname.sh
```

Notice that the hostname here must include the domain. So, our hostname was *pinger.raspberry.slac.stanford.edu*.

Then, we followed the instructions in [PingER End-to-end Reporting version 2](#). After installing the [PingER2](#) monitoring code, we installed the ping_data gathering agent, the [traceroute server](#) and the pinger_trimmer following the [instructions](#).

When we tested everything, we got a error message on the *pingerCronStat.stderr* file telling us that the *ParserDetails.ini* file was missing. We used [this approach](#) to fix this.

We entered the machines as monitors in the PingER meta data base of hosts.

Obs: Make sure to change the default password for Raspbian.

Requirements

Two major points need to be addressed before we can comfortably deploy Raspberry PI MAs.

1. The Raspberry Pi PingER MA must be robust and reliable. It needs to run for months to years with no need for intervention and needs automatic recovery after power outages.
2. The important metrics derived from the measurements made by the Raspberry Pi should not be significantly different from those made by a bare metal PingER MA, or if they are then this needs to be understood.
3. For applications where there is no reliable power source, the Raspberry Pi needs to be able to run 24 hours a day with only solar derived power. Let's say the power required is 3W at 5V or $(3/5)A=0.6A$. If we have a 10,000 mAh battery, then at 0.6A it should have power for $10Ah/0.6A \sim 16.6$ hours. Then we need a solar cell to be able to refill the battery in a few hours of sunlight. Let's take a 20W 5 V solar panel (see <http://www.amazon.com/SUNKINGDOM-trade-Foldable-Charger-Compatible/dp/B00MTEDTWG>) = $20/5 = 4$ A solar panel. So initial guess is $10A-h/4A = 2.5$ hours. But inefficiencies (see <http://www.voltaicsystems.com/blog/estimating-battery-charge-time-from-solar/>) of say 2.5 extends this to 6.25 hours.

We define the important metrics measured by PingER as being the minimum, average, median and jitter of the RTTs, the packet loss, together with the reachability (i.e. a target host is unreachable when no ping requests are responded by the target host). These are the main metrics that impact applications such as throughput, voice over IP, streaming video, haptics, and estimating the [geolocation of a host by pinging it from well know landmarks](#). Such differences might result in significant discontinuities in the metric measurements if we were to change the monitoring host from a bare metal server to a Raspberry Pi.

In this report the jitter is represented by the [Inter Packet Delay \(IPD\)](#), the absolute values of the IPF (Abs(IPD)) and the Stdev(IPD).

Methodology

The measurement tool was a PingER MA in all cases. PingER (Ping End-to-end Reporting) is the name given to the Internet End-to-end Performance Measurement (IEPM) project to monitor end-to-end performance of Internet links. Originally, in 1995 it was for the High Energy Physics community, however, this century it has been more focused on measuring the Digital Divide from an Internet Performance viewpoint. The project now involves measurements to **over 700 sites** in over 160 countries. It uses the ubiquitous [ping facility](#) so no special software has to be installed on the targets being measured by the MAs.

PingER measurements are made by ~60 MAs in 23 countries. They make measurements to over 700 targets in ~ 160 countries containing more than 99% of the world's connected population. The measurement cycle is scheduled at roughly 30 minute intervals. The actual scheduled timing of a measurement is deliberately randomized so measurements from one MA are not synchronized with another MA. Typical absolute separation of the timestamp of a measurement from say `pinger.slac.stanford.edu` to `sitka.triumf.ca` versus `pinger-raspberry.slac.stanford.edu` to `sitka.triumf.ca` is several minutes (e.g. ~ 8 mins for measurements during the time frame June 17 to July 14, 2015), see spreadsheet. At each measurement cycle, each MA issues a set of 100 Byte pings and a set of 1000 Byte ping requests to each target in the MA's list of targets, stopping when the MA receives 10 ping responses or it has issued 30 ping requests. The number of ping responses is referred to as N and is in the range 0 - 10. The data recorded for each set of pings consists of: the MA and target names and IP addresses; a time-stamp; the number of bytes in the ping request; the number of ping requests and responses (N); the minimum Round Trip Time (RTT) (Min_RTT), the average RTT (Avg_RTT) and maximum RTT (Max_RTT) of the N ping responses; followed by the N ping sequence numbers, followed by the N RTTs. From the N RTTs we derive various metrics including: the minimum ping RTT; average RTT; maximum RTT; standard deviation (stdev) of RTTs, 25% probability (first quartile) of RTT; 75% probability (third quartile) of RTT; Inter Quartile Range (IQR); loss; and reachability (host is unreachable if it gets 100 % loss). We also derive the Inter Packet delay (IPD) and the Inter Packet Delay Variability (IPDV) as the IQR of the IPDs.

The data is publicly available and since the online data goes back to January 1998, it currently provides ~19 years of historical data on worldwide Internet performance.

Measurements

We chose to make detailed measurements to and from two MAs at SLAC.

1. The [Dell Poweredge 2650](#) bare metal `pinger.slac.stanford.edu` server running Red Hat Linux 2.6.32-504.8.1.el6.i686 [with specs](#).

```
103cottrell@pinger:~$uname -a
Linux pinger 2.6.32-504.8.1.el6.i686 #1 SMP Fri Dec 19 12:14:17 EST 2014 i686 i686 i386 GNU/Linux
```

2. The Raspberry Pi `pinger-raspberry.slac.stanford.edu` an armv61 running Gnu Linux (see above) [with specs](#).

```
pi@pinger-raspberry ~ $ uname -a
Linux pinger-raspberry.slac.stanford.edu 3.18.11+ #781 PREEMPT Tue Apr 21 18:02:18 BST 2015 armv61 GNU
/Linux
```

Both were in the same building at SLAC, i.e. roughly at latitude 37.4190 N, longitude 122.2085 W, but on different floors. The machines are about 30 metres apart or about 0.0003 msec based on the speed of light in a direct fibre.

The measurements were made:

- between `pinger.slac.stanford.edu` and `pinger-raspberry.slac.stanford.edu`
- from both `pinger.slac.stanford.edu` and `pinger-raspberry.slac.stanford.edu` to targets at varying distances and hence varying minimum RTTs from SLAC.
- from two representative MAs at `sitka.triumf.ca` and `ping.cern.ch` to `pinger.slac.stanford.edu` and `pinger-raspberry.edu`

Information on the various hosts involved is given in the table below.

The Directivity in the table below provides a measure of how direct the route is between the MA and target. The Directivity is given as:

*Directivity = great circle distance between MA & target [in km] / (RTT [ms] * 100 [km/ms])*

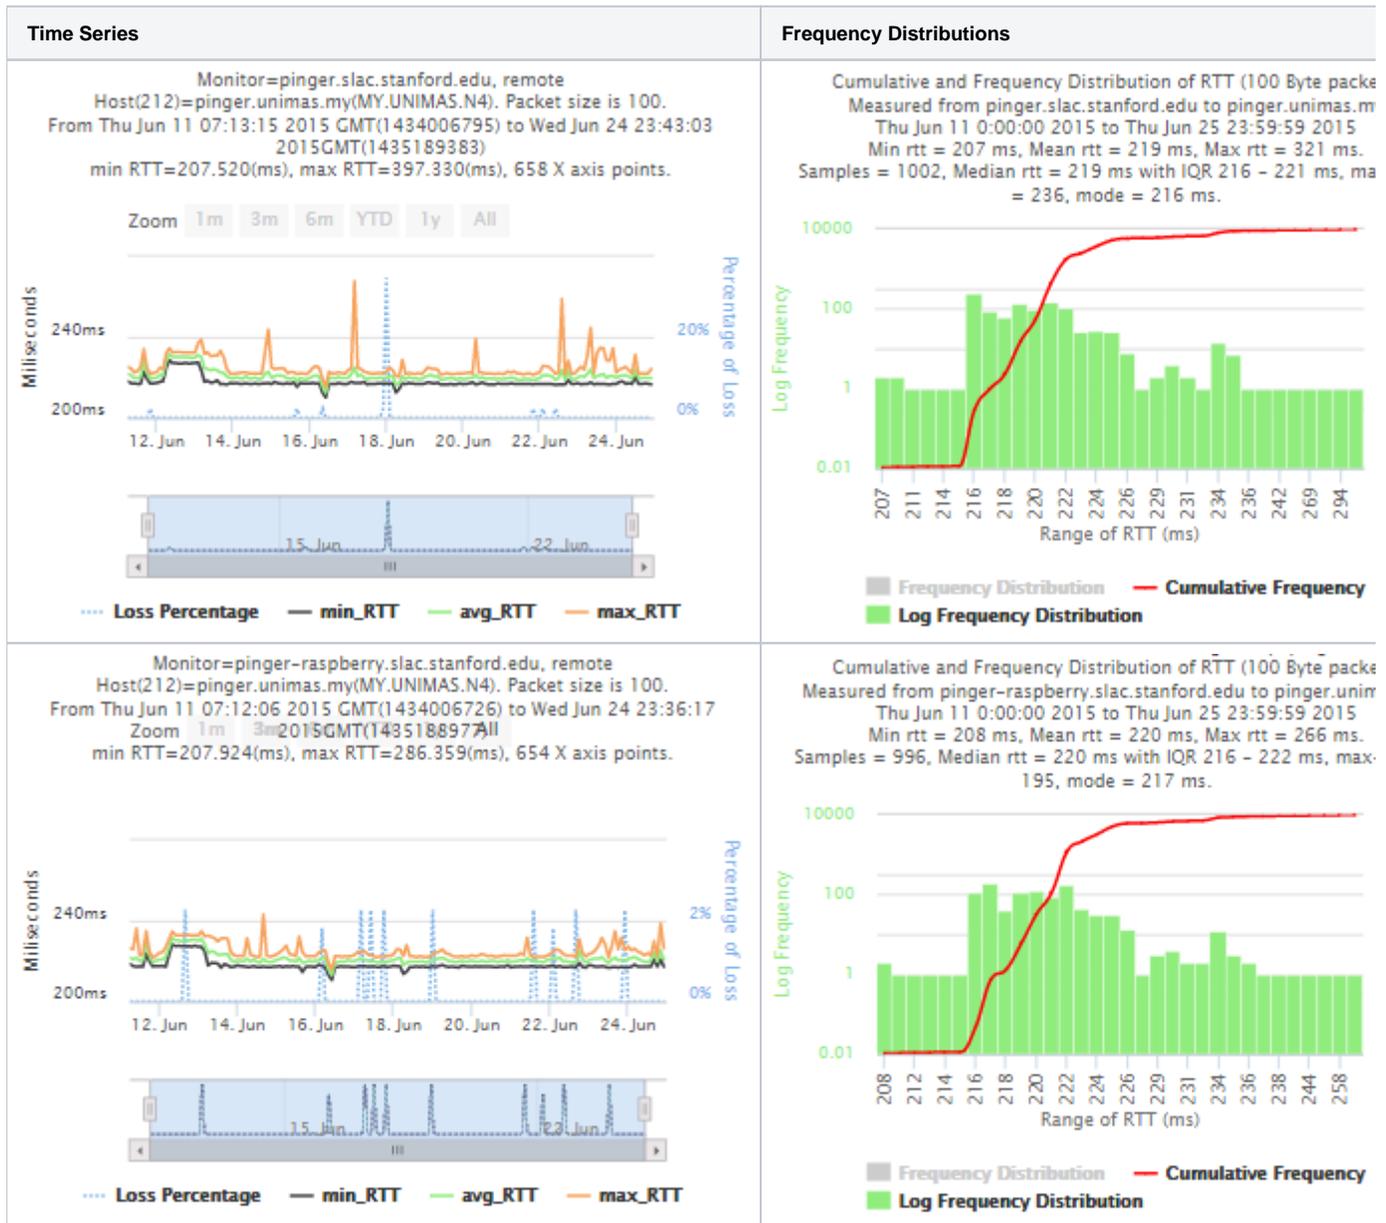
The Directivity is ≤ 1 , and a value of 1 means the RTT is the same as given by the speed of light in a fibre.

Host	Lat	Long	Great Circle distance from SLAC	Min RTT (as constrained by speed of light in fibre)	Directivity based on measured min RTT
<code>pinger.slac.stanford.edu</code>	37.4190 N	122.2085 W	0 km	0.0003 ms	0.001
<code>pinger-raspberry.slac.stanford.edu</code>	37.4190 N	122.2085 W	0 km	0.0003 ms	0.001
<code>sitka.triumf.ca</code>	49.2475 N	123.2308 W	1319.6 km	13.196 ms	0.6
<code>ping.cern.ch</code>	46.23 N	6.07 E	9390.6 km	93.90 ms	0.63

The following plots are the raw metric data from the MA measurements from which we have to craft our text (with some plot examples).

Example target = pinger.unimas.my (~220 msec.)

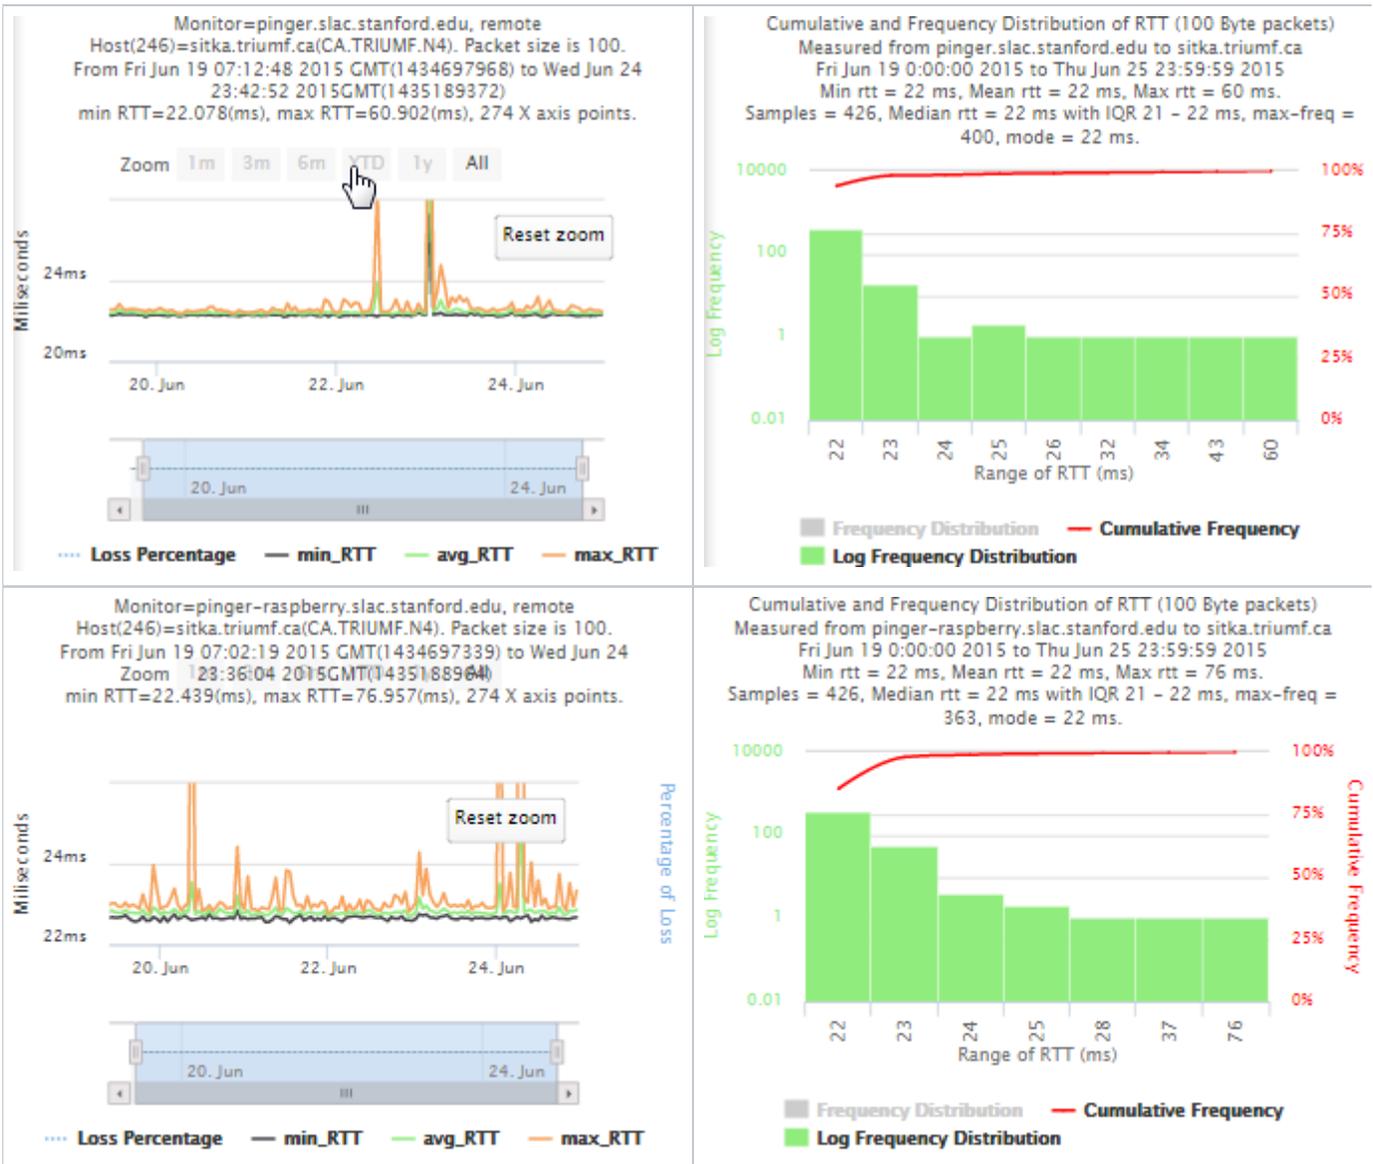
For both 100Byte and 1000 Byte pings (not shown above) the round trip time series for RTTs have similar behaviour and there are similar losses 7:10 (pinger : pinger-raspberry for 100 Byte pings), note the different Y scales for losses. The losses are about double for 1000Byte pings.



Example target sitka.triumf.ca (~22msec.)

For 100Byte the round trip time series for RTTs did not have similar behaviour. We noticed a great change mainly in the maximum round trip time. The average minimum RTT did not change that much. Another point about pinger-raspberry is that it increases significantly the RTT for near nodes (about ~1ms). The difference is greater than if we compare a node which is in a long distance.

Time series	Frequency distributions
-------------	-------------------------



pinger.slac.stanford.edu and pinger-raspberry.slac.stanford.edu

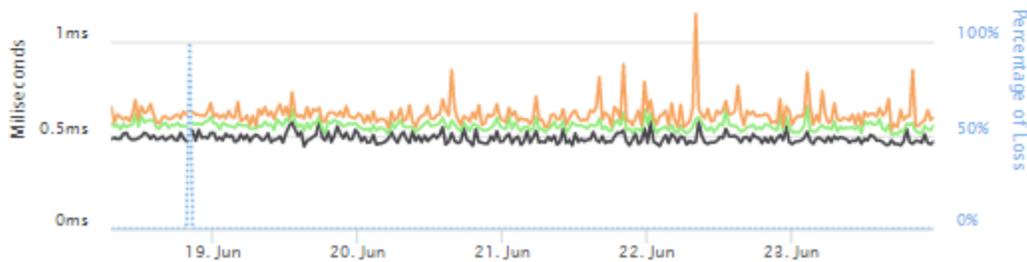
Now, we compared the RTT between pinger and pinger-raspberry. They are located in the same network and the RTT should be very small. However, as noticed before pinger-raspberry has a greater maximum RTT than pinger. The average RTT also has some difference, but now as much as the maximum time has. Note that the second graph represents the third graph using the same scale as the first (pinger graph).

Monitor=pinger.slac.stanford.edu, remote Host(184) is pinger-raspberry.slac.stanford.edu(EDU.SLAC.STANFORD.RASPBERRY). Packet size is 100.
From Thu Jun 18 07:07:20 2015 GMT(1434611240) to Tue Jun 23 23:37:32 2015 GMT(1435102652)
min RTT=0.441(ms), max RTT=1.152(ms), 274 X axis points.
[\[Form\]](#)[\[Add or Remove csv data\]](#) [\[Frequency Distribution\]](#) [\[Get csv data for Excel\]](#) [\[Help\]](#)

Comparison of minimum, average and maximum roundtrip times, plus losses ☰

Source: www-wanmon.slac.stanford.edu/cgi-wrap/pingtable.pl

Zoom From To

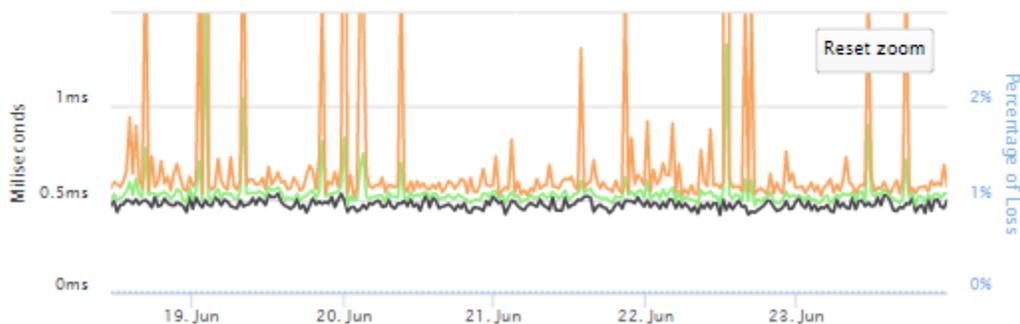


Monitor=pinger-raspberry.slac.stanford.edu, remote Host(207) is pinger.slac.stanford.edu(EDU.SLAC.STANFORD.N3). Packet size is 100.
From Thu Jun 18 07:19:04 2015 GMT(1434611944) to Tue Jun 23 23:56:37 2015 GMT(1435103797)
min RTT=0.418(ms), max RTT=10.795(ms), 272 X axis points.
[\[Form\]](#)[\[Add or Remove csv data\]](#) [\[Frequency Distribution\]](#) [\[Get csv data for Excel\]](#) [\[Help\]](#)

Comparison of minimum, average and maximum roundtrip times, plus losses ☰

Source: www-wanmon.slac.stanford.edu/cgi-wrap/pingtable.pl

Zoom From To

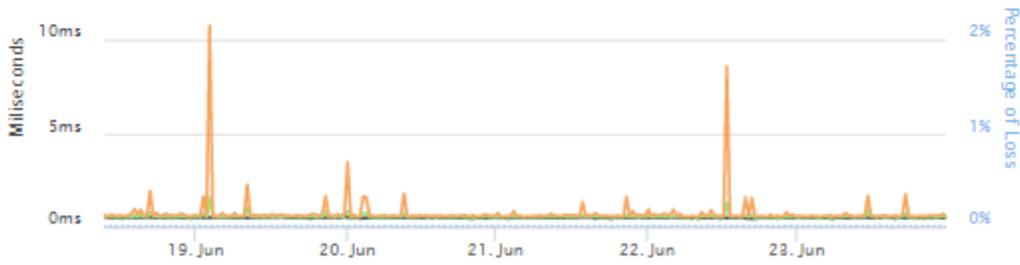


Monitor=pinger-raspberry.slac.stanford.edu, remote Host(207) is
pinger.slac.stanford.edu(EDU.SLAC.STANFORD.N3). Packet size is 100.
From Thu Jun 18 07:19:04 2015 GMT(1434611944) to Tue Jun 23 23:56:37 2015 GMT(1435103797)
min RTT=0.418(ms), max RTT=10.795(ms), 272 X axis points.
[\[Form\]](#)[\[Add or Remove csv data\]](#) [\[Frequency Distribution\]](#) [\[Get csv data for Excel\]](#) [\[Help\]](#)

Comparison of minimum, average and maximum roundtrip times, plus losses ☰

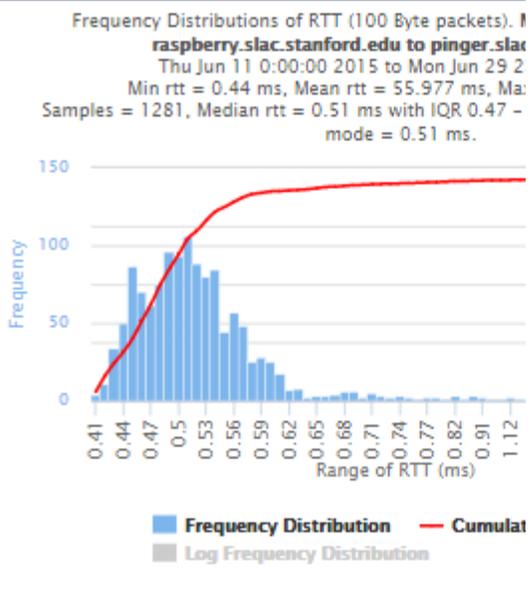
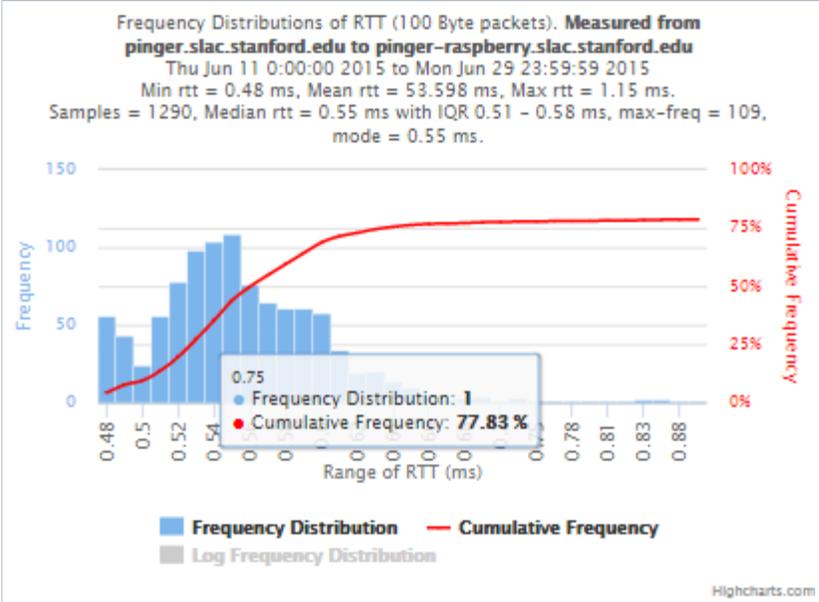
Source: www-wanmon.slac.stanford.edu/cgi-wrap/pingtable.pl

Zoom From To



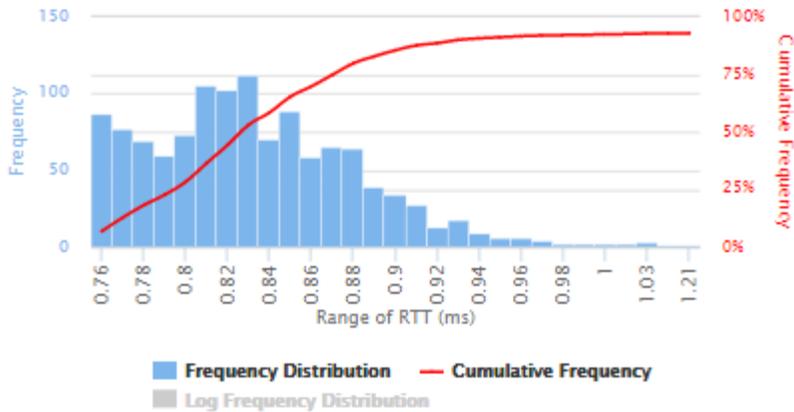
pinger to Pinger-raspberry

pinger-raspberry to pinger



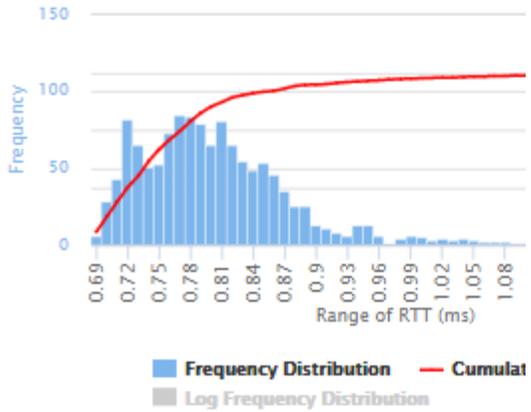
PingER Data Frequency Plots

Source: www-wanmon.slac.stanford.edu/cgi-wrap/pingtable.pl
 Frequency Distributions of RTT (1000 Byte packets). **Measured from pinger.slac.stanford.edu to pinger-raspberry.slac.stanford.edu**
 Thu Jun 11 0:00:00 2015 to Mon Jun 29 23:59:59 2015
 Min rtt = 0.76 ms, Mean rtt = 82.671 ms, Max rtt = 1.21 ms.
 Samples = 1290, Median rtt = 0.83 ms with IQR 0.79 - 0.86 ms, max-freq = 112, mode = 0.83 ms.



PingER Data Frequency

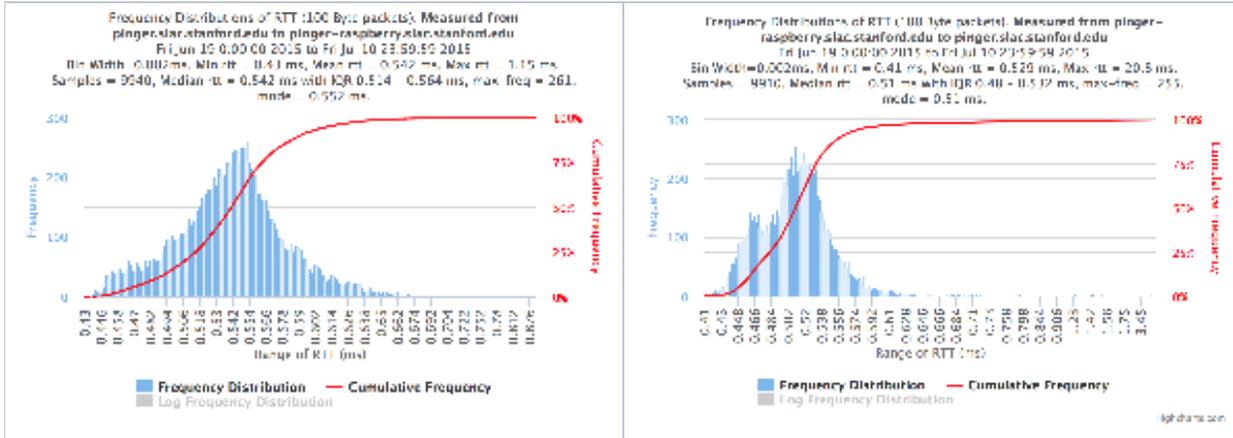
Source: www-wanmon.slac.stanford.edu/cgi-wrap/pingtable.pl
 Frequency Distributions of RTT (1000 Byte packets). **Measured from raspberry.slac.stanford.edu to pinger.slac**
 Thu Jun 11 0:00:00 2015 to Mon Jun 29 2 00:00:00 2015
 Min rtt = 0.76 ms, Mean rtt = 81.147 ms, Max rtt = 1.08 ms.
 Samples = 1281, Median rtt = 0.82 ms with IQR 0.77 - 0.86 ms, max-freq = 112, mode = 0.77 ms.



Using full set of pings for RTT frequency distributions

The frequency plots above are for the frequencies of the minimum, average and maximum RTTs. Below we show the frequencies when we take the individual pings (usually 10 assuming little loss) for all the ping RTTs in each measurement set.

pinger to pinger-raspberry	pinger-raspberry to pinger
----------------------------	----------------------------



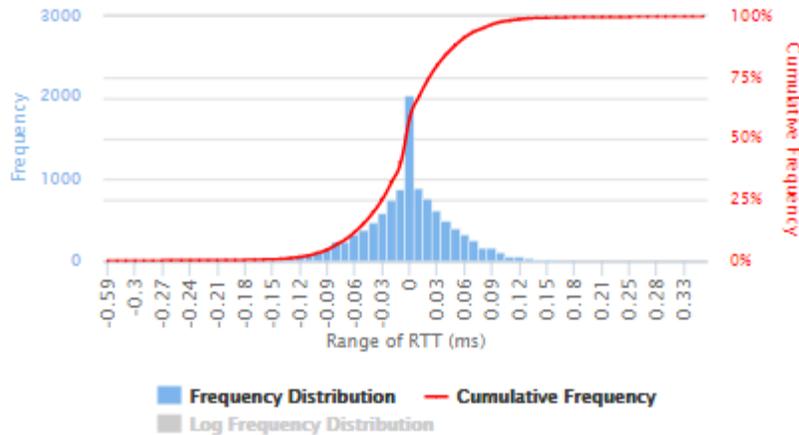
Frequency distribution for absolute interpacket delays.

The magnitude of the RTT is very dependent on the distance of the path between the source and destination. Many applications such as voice over IP, video streaming, or haptics are very dependent on the variability or jitter of the RTT. The jitter is often more dependent on the network edges compared to the RTT. There are many ways to calculate the jitter (see for example <http://www.slac.stanford.edu/comp/net/wan-mon/tutorial.html#variable>). We calculate the inter packet delay (IPD) and the absolute IPD and display the frequency distributions and statistics.

pinger to pinger-raspberry	pinger-raspberry to pinger
-----------------------------------	-----------------------------------

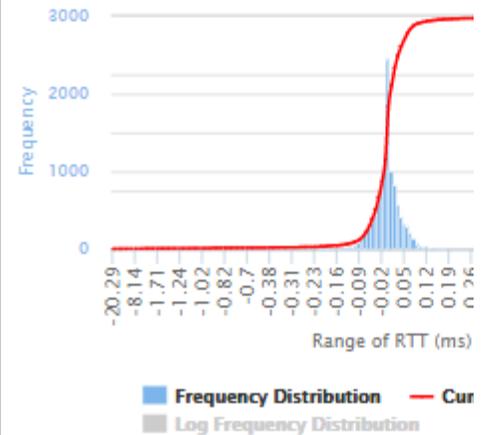
IPD

Frequency Distributions of ipd RTT (100 Byte packets). Measured from **pinger.slac.stanford.edu to pinger-raspberry.slac.stanford.edu**
 Wed Jun 17 0:00:00 2015 to Mon Jul 13 23:59:59 2015
 Bin Width=0.01 ms, Min = -0.59 ms, Mean = -0.000 ms, Max = 0.62 ms.
 Samples = 10890, Median = 0 ms with IQR (0.03 - -0.04) = 0.070 ms, max-freq = 2035, mode = 0 ms.

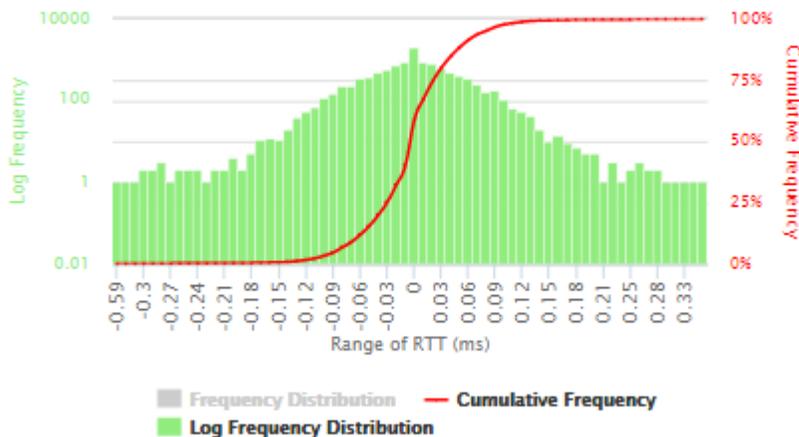


Highcharts.com

Frequency Distributions of ipd RTT (100 Byte packets). Measured from **pinger-raspberry.slac.stanford.edu to pinger.slac.stanford.edu**
 Wed Jun 17 0:00:00 2015 to Mon Jul 13 23:59:59 2015
 Bin Width=0.01 ms, Min = -20.29 ms, Mean = -0.000 ms, Max = 0.26 ms.
 Samples = 10850, Median = 0 ms with IQR (0.02 - -0.02) = 0.040 ms, mode = 0 ms.

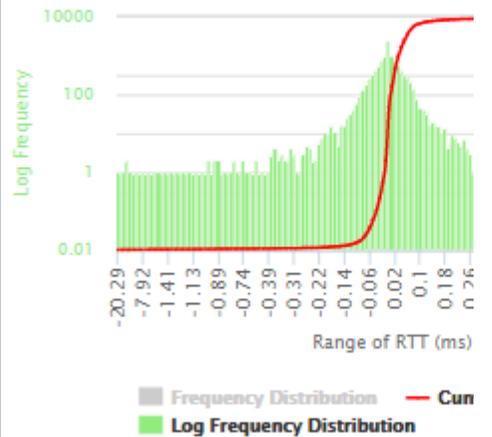


Frequency Distributions of ipd RTT (100 Byte packets). Measured from **pinger.slac.stanford.edu to pinger-raspberry.slac.stanford.edu**
 Wed Jun 17 0:00:00 2015 to Mon Jul 13 23:59:59 2015
 Bin Width=0.01 ms, Min = -0.59 ms, Mean = -0.000 ms, Max = 0.62 ms.
 Samples = 10890, Median = 0 ms with IQR (0.03 - -0.04) = 0.070 ms, max-freq = 2035, mode = 0 ms.



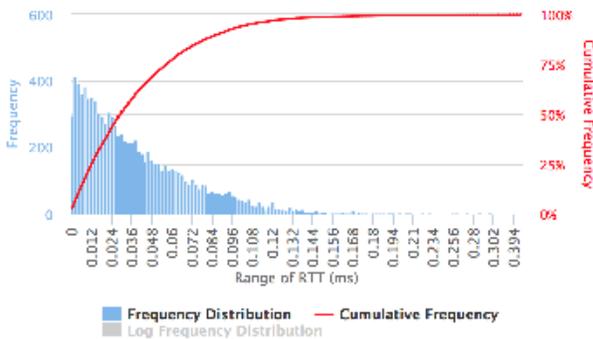
Highcharts.com

Frequency Distributions of ipd RTT (100 Byte packets). Measured from **pinger-raspberry.slac.stanford.edu to pinger.slac.stanford.edu**
 Wed Jun 17 0:00:00 2015 to Mon Jul 13 23:59:59 2015
 Bin Width=0.01 ms, Min = -20.29 ms, Mean = -0.000 ms, Max = 0.26 ms.
 Samples = 10850, Median = 0 ms with IQR (0.02 - -0.02) = 0.040 ms, mode = 0 ms.

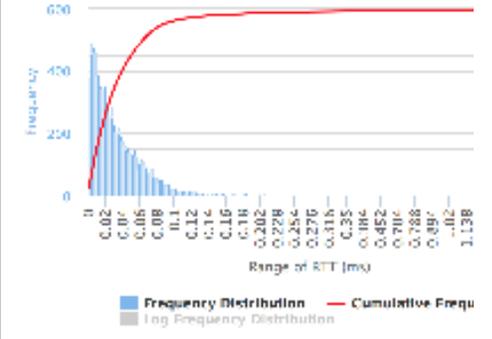


Abs (IP D)

Frequency Distributions of abs(ipd) RTT (100 Byte packets). Measured from **pinger.slac.stanford.edu to pinger-raspberry.slac.stanford.edu**
 Fri Jun 19 0:00:00 2015 to Fri Jul 10 23:59:59 2015
 Bin Width=0.002ms, Min rtt = 0 ms, Mean rtt = 0.041 ms, Max rtt = 0.628 ms.
 Samples = 9940, Median rtt = 0.03 ms with IQR 0.01 - 0.058 ms, max-freq = 414, mode = 0.002 ms.



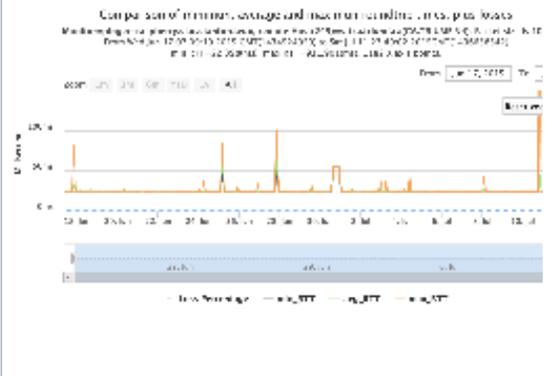
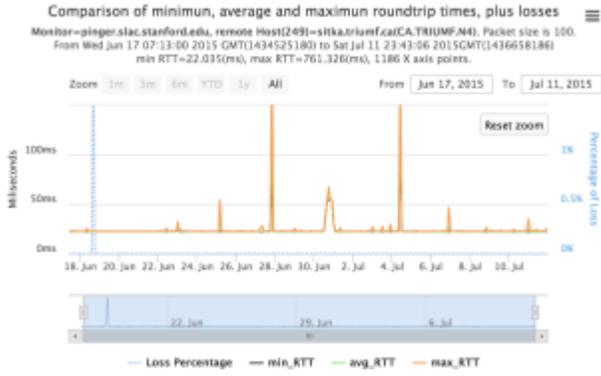
Frequency Distributions of abs(ipd) RTT (100 Byte packets). Measured from **pinger-raspberry.slac.stanford.edu to pinger.slac.stanford.edu**
 Fri Jun 19 0:00:00 2015 to Fri Jul 10 23:59:59 2015
 Bin Width=0.002ms, Min rtt = 0 ms, Mean rtt = 0.056 ms, Max rtt = 0.628 ms.
 Samples = 9910, Median rtt = 0.024 ms with IQR 0.006 - 0.051 ms, max-freq = 490, mode = 0.002 ms.



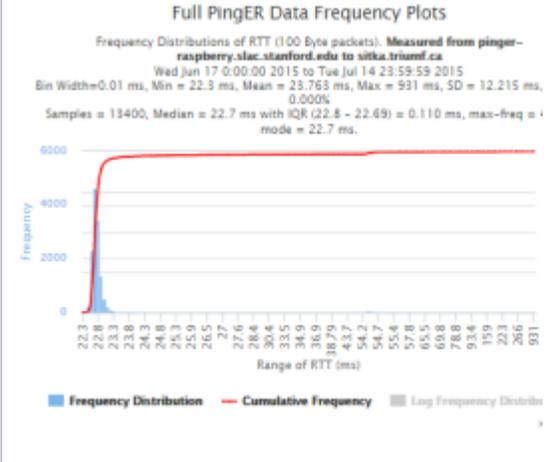
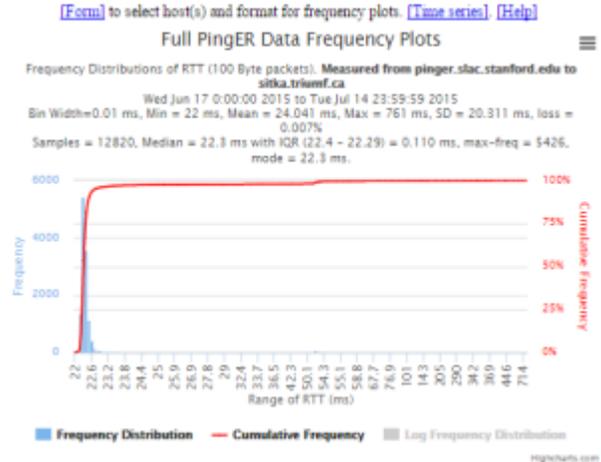
pinger.slac.stanford.edu to sitka.triumf.ca

pinger-raspberry.slac.stanford.edu to sitka.triumf.ca

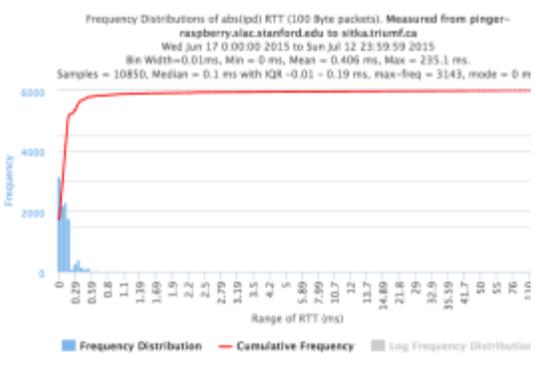
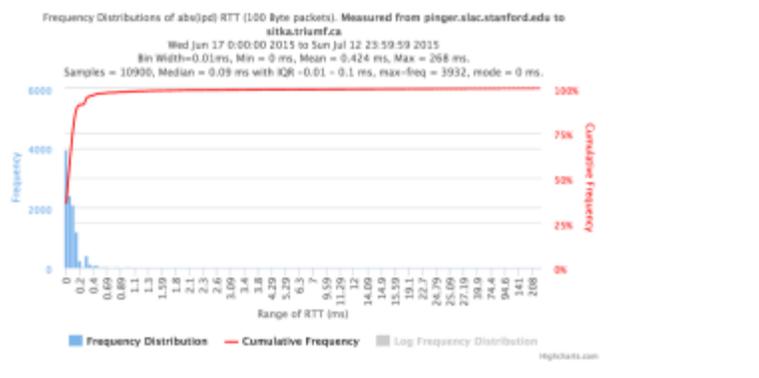
Time Series



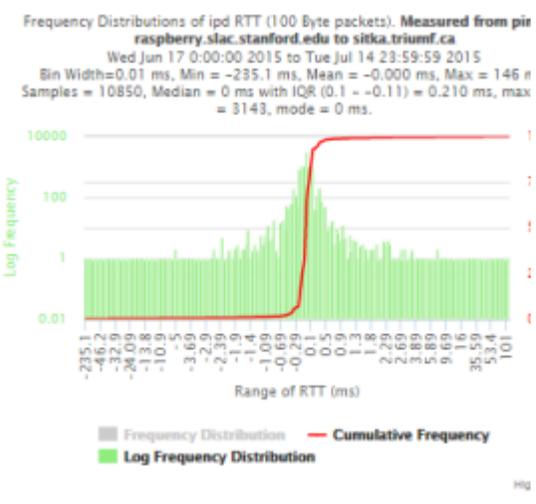
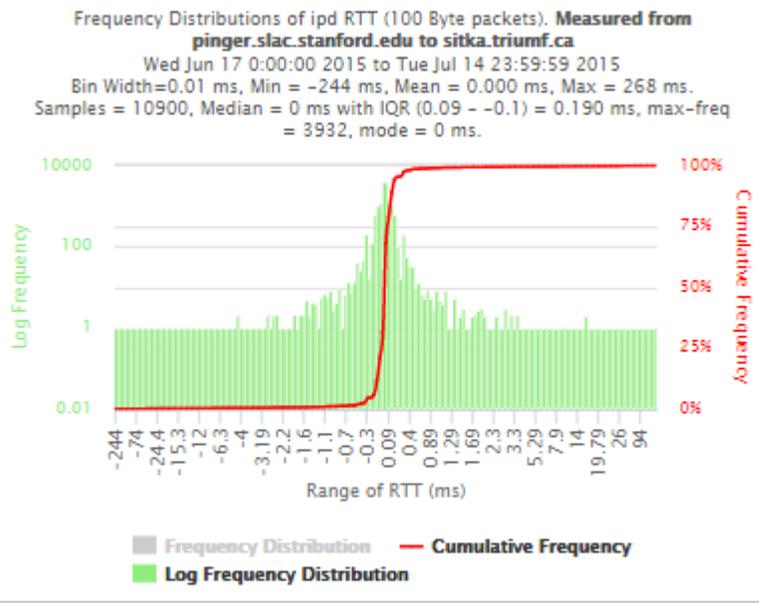
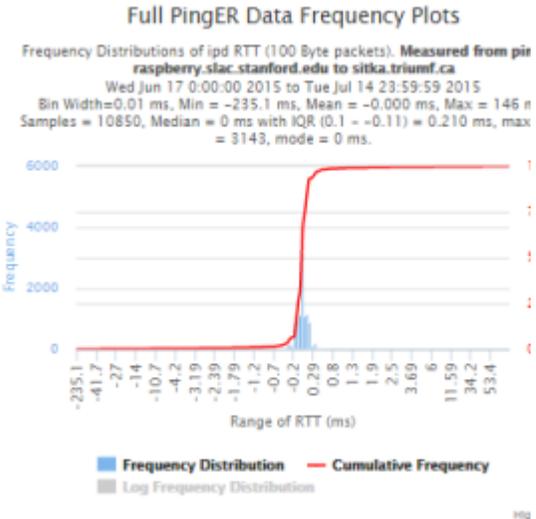
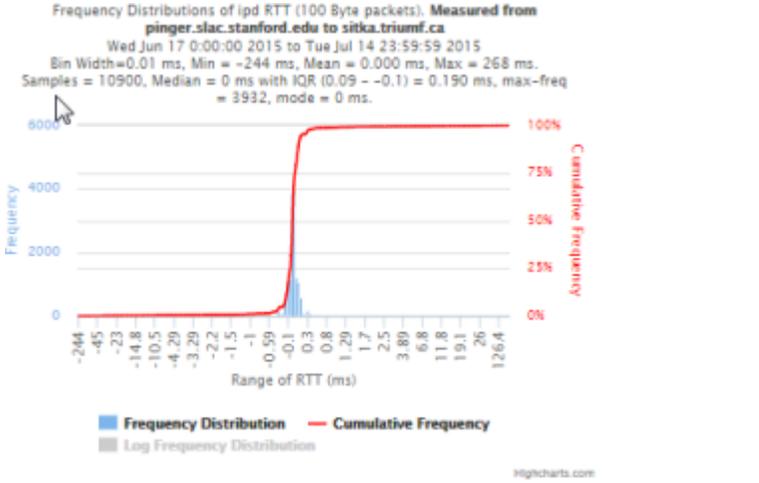
Frequency distribution



Frequency distribution Abs (IP D)



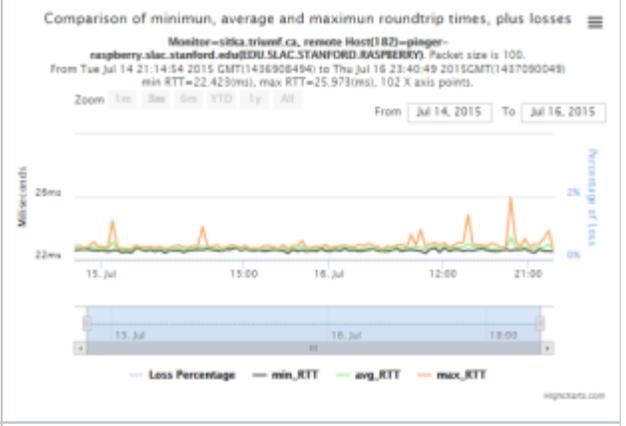
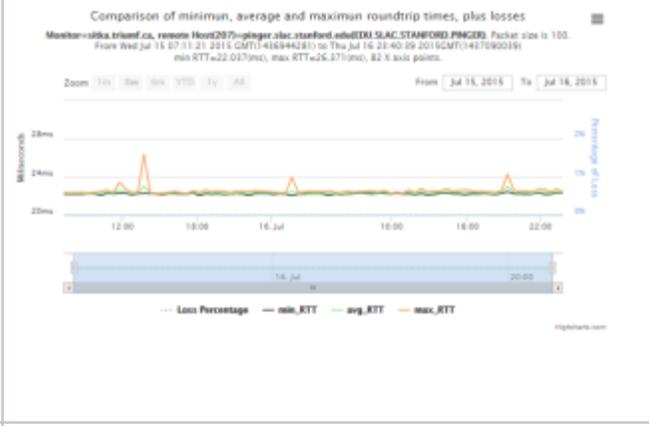
Frequency distribution IPD



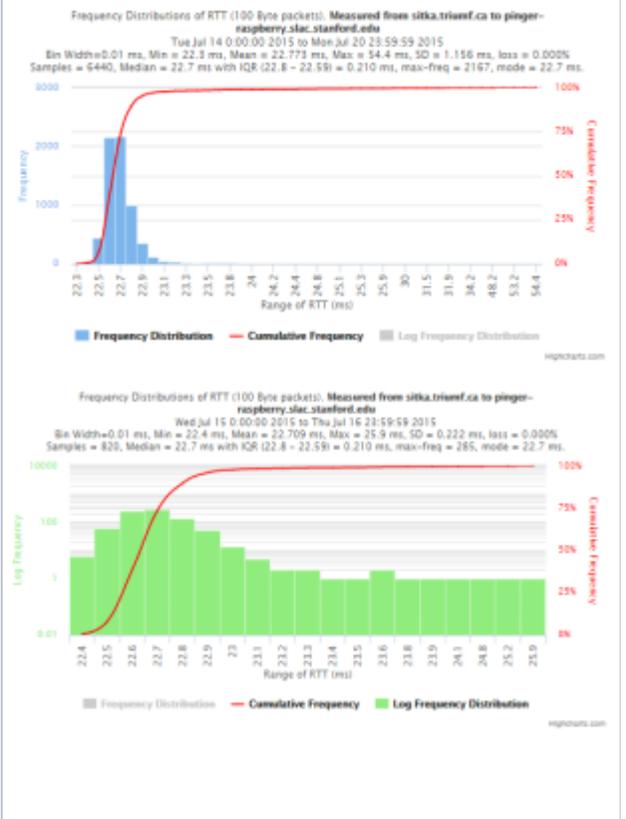
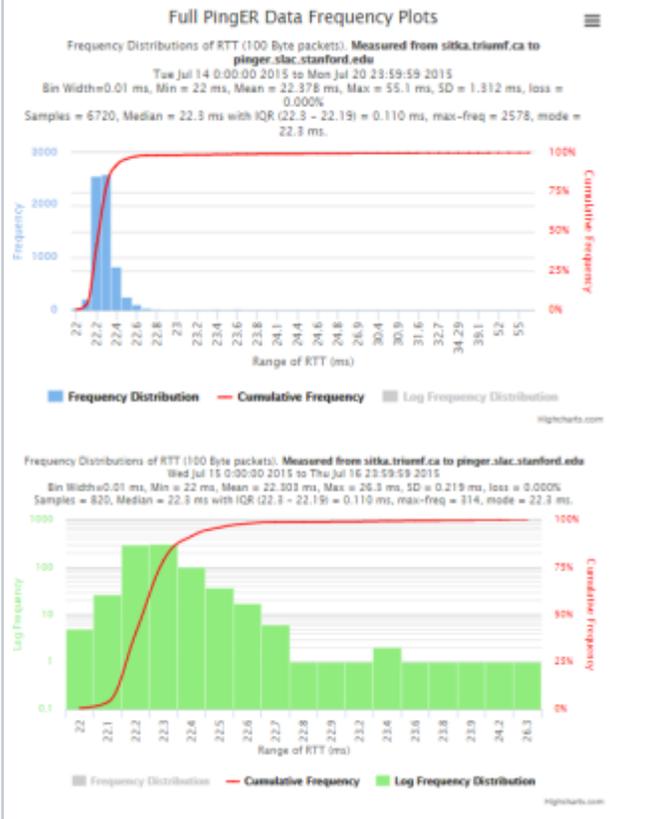
From sitka.triumf.ca to SLAC

	sitka.triumf.ca to pinger.slac.stanford.edu	sitka.triumf.ca to pinger-raspberry.slac.stanford.edu
--	---	---

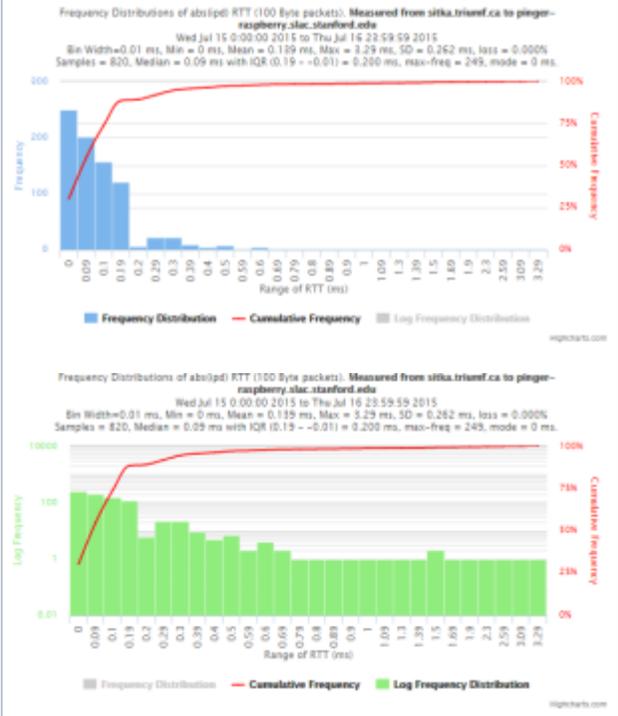
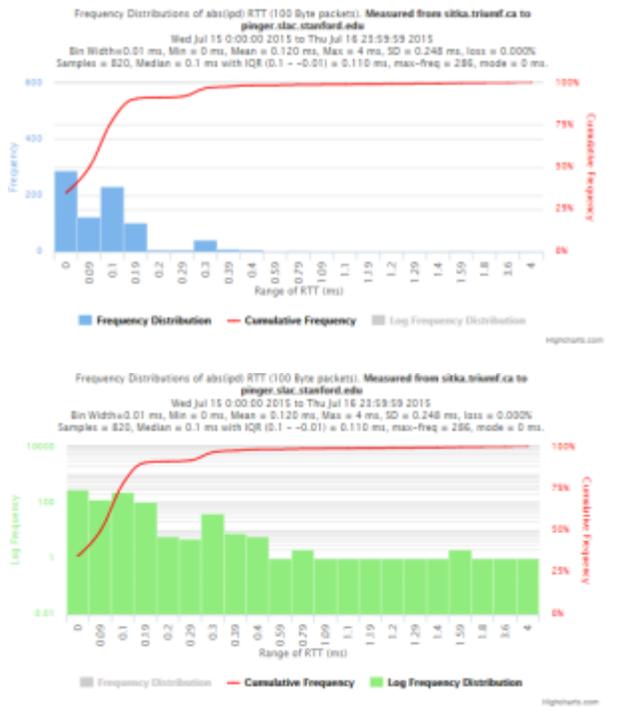
Time series



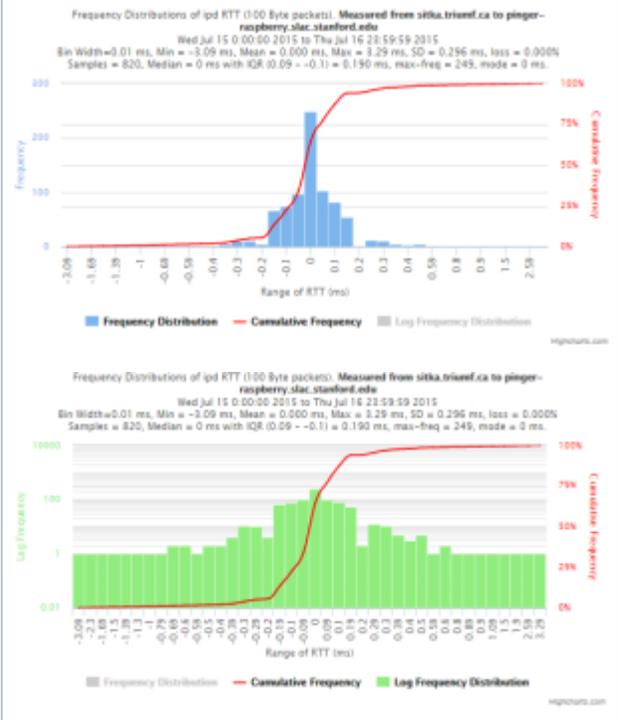
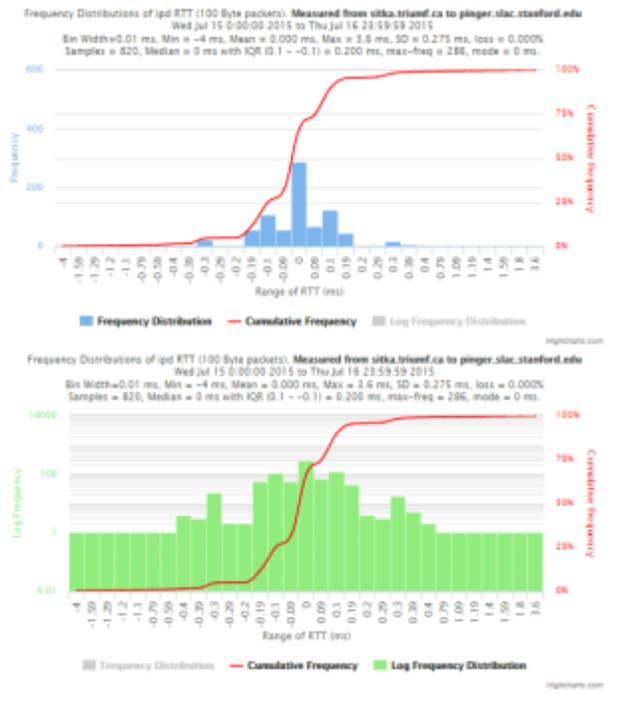
Frequency Distribution of RTT



Frequency distribution of Abs (IPD)

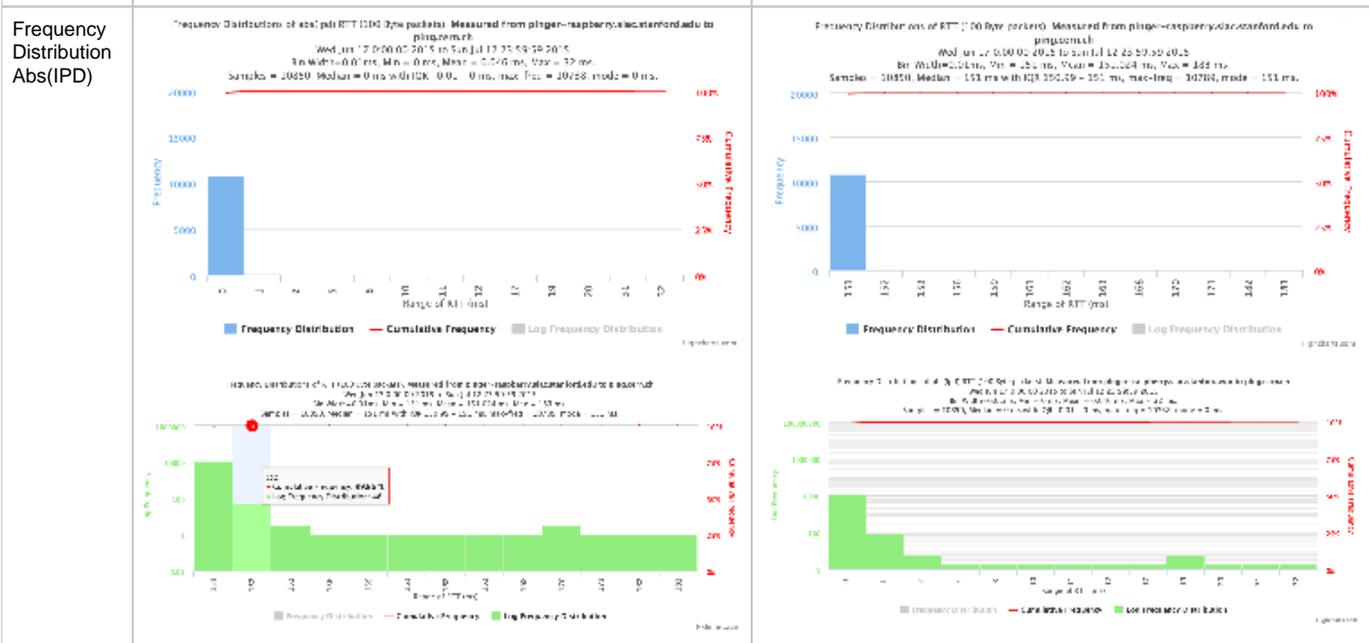
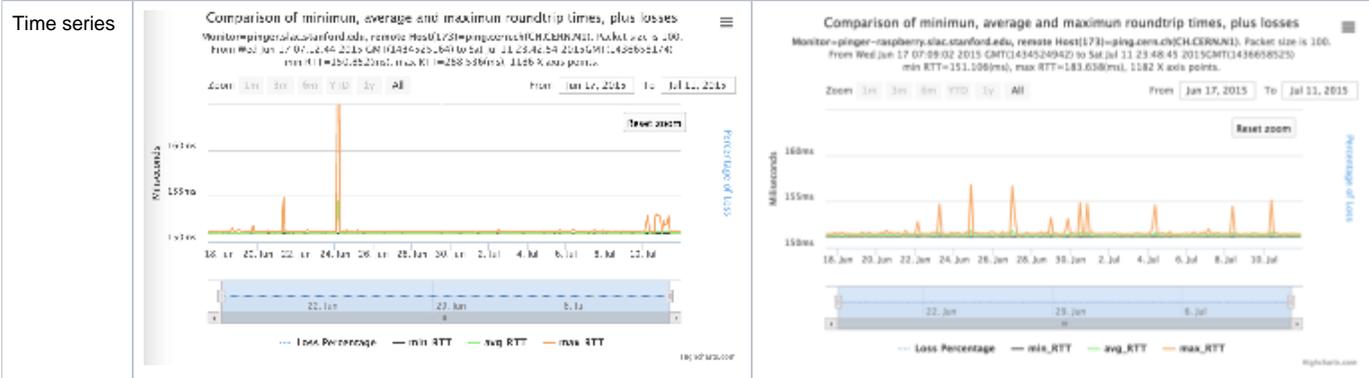


Frequency distribution of IPD

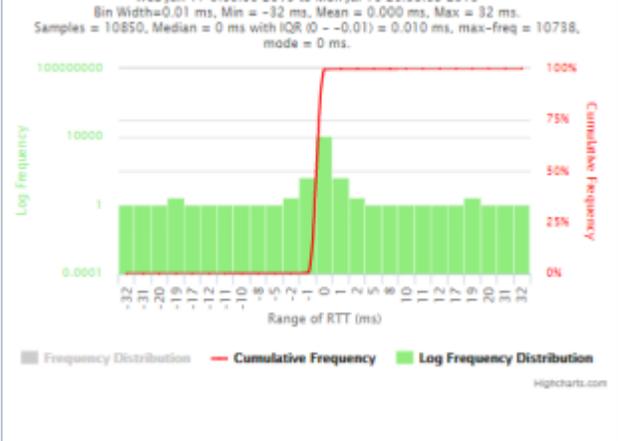
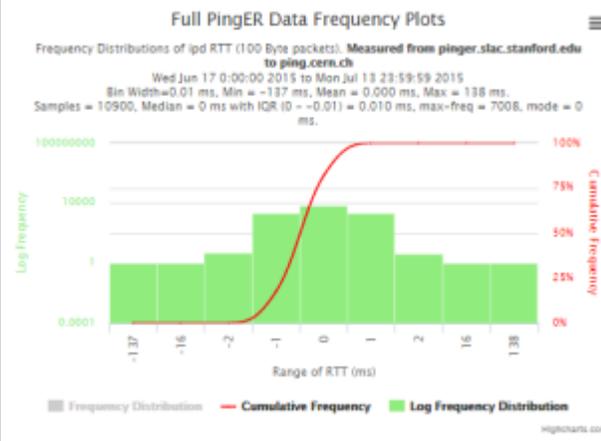
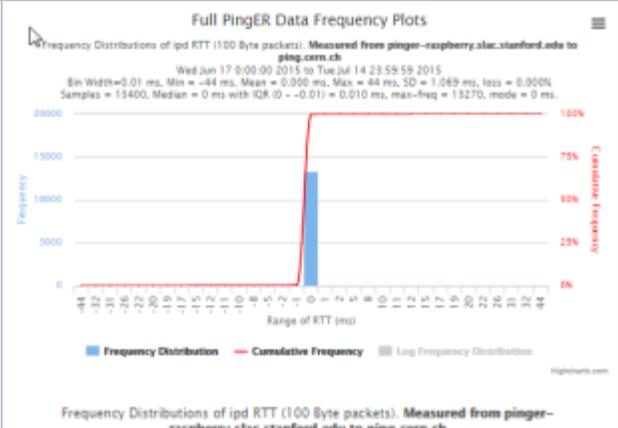
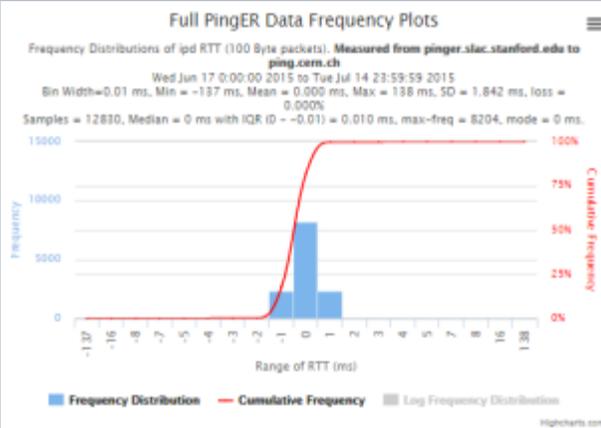


To ping.cern.ch from SLAC

To ping.cern.ch from pinger	To ping.cern.ch from pinger-raspberry
-----------------------------	---------------------------------------



Frequency Distribution IPD



Results

The table below shows the various aggregate metrics measured from MA to target comparing pinger.slac.stanford with pinger-raspberry.slac.stanford.edu. The columns are arranged in pairs. The first of each pair is for pinger.slac.stanford.edu, the second for pinger-raspberry.slac.stanford.edu. Each pair is measured over the same time period. Different pairs are measured over different time periods.

The errors (+/-) are estimated as follows:

- The Min RTT error is the IQR of the minimum RTTs from each set of pings for the time period.
- The Avg RTT error is the standard deviation of all the average RTTs from each set of pings for the time period
- The Median RTT error is the IQR of all the individual RTTs from all sets of measurements for the time period
- The IPDV error is the IQR of the hourly IQRs for the time period (see [Spreadsheet](#))

Metric \ Monitor Target	pinger to pinger-raspberry	pinger-raspberry to pinger	pinger to sitka	pinger-raspberry to sitka	sitka to pinger	sitka to pinger-raspberry	pinger to CERN	pinger-raspberry to CERN
Time period	June 17 - July 13, 2015	June 17 - July 13, 2015	June 17 - July 14, 2015	June 17 - July 14	July 15 - July 16 2015	July 15 - July 16 2015	June 17 - July 12, 2015	June 17 - July 12, 2015
Samples	10890	10850	10900	10850	820	820	10900	10900
Min RTT	0.43 ms	0.41 ms	22 +- 1 ms	22.3 +- 1 ms	22 ms	22.4 ms	150 +- 1 ms	151 ms
Avg RTT	0.542 +- 0.06 ms	0.529 +- 0.5ms	23.9 +- 1ms	23.827 +- 1 ms	22.30 ms	22.709 ms	150.307 +- 2.4 ms	151.024 +- 1.3 ms
Max RTT	1.15 ms	20.8 ms	761 ms	334 ms	26.3 ms	25.9 ms	288 ms	183 ms
Stdev	0.055 ms	0.540 ms	9.58 ms	9.60 ms	0.219 ms	0.222 ms	2.37 ms	1.3 ms
Median RTT	0.542 +- 0.05 ms	0.51 ms +- 0.052 ms	22.3 +- 0.2 ms	22.7 +- 0.11 ms	22.3 +- 0.11 ms	22.7 +- 0.21 ms	150 +- 1 ms	151 +- 0.01 ms
25%	0.514 ms	0.48 ms	22.2 ms	22.69 ms	22.3 ms	22.8 ms	149.99 ms	150.99 ms
75%	0.564 ms	0.532 ms	22.4 ms	22.8 ms	22.19 ms	22.59 ms	151 ms	151 ms
IQR	0.05 ms	0.052 ms	0.2 ms	0.11 ms	0.113 ms	0.210 ms	1.01 ms	0.01 ms
Min IPD	-0.59 ms	-20.29 ms	-244 ms	-235.1 ms	-0.4 ms	-0.39 ms	-137 ms	-32 ms
Avg IPD	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms	0
Max IPD	0.62 ms	20.26 ms	268 ms	146 ms	3.6 ms	3.29 ms	138 ms	32 ms

Median IPD	0 +- 0.07 ms	0 +- 0.05 ms	0 +- 0.2 ms	0 +- 0.2ms	0 ms +- 0.2 ms	0 +- 0.2 ms	0 +- 0.01 ms	0 +- 0.01 ms
Stdev					0.248 ms	0.296 ms		
25% IPD	-0.04 ms	-0.03 ms	-0.1 ms	-0.11 ms	-0.1 ms	-0.1 ms	-0.01 ms	-0.01 ms
75% IPD	0.03 ms	0.02 ms	0.09 ms	0.1 ms	0.01 ms	0.09 ms	0 ms	0 ms
IQR IPD	0.07 ms	0.05 ms	0.190 +- 0.5 ms	0.210 +- 0.4 ms	0.2 ms	0.190ms	0.01 ms	0.01 ms
Min(abs(IPD))	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms
Avg(abs(IPD))	0.041 ms	0.066 ms	0.424 ms	0.406 ms	0.120 ms	0.139 ms	0.386 ms	0.046 ms
Max(abs(IPD))	0.0628ms	20.294 ms	268 ms	235.1 ms	4 ms	3.29 ms	138 ms	32 ms
Stdev					0.248 ms	0.262 ms		
Median(abs(IPD))	0.03ms	0.024 ms	0.09 ms	0.1 ms	0.1 ms	0.09 ms	0 ms	0 ms
25%(abs(IPD))	0.01ms	0.008 ms	0.01 ms	0.01 ms	-0.01 ms	-0.01 ms	0.01 ms	0.01 ms
75%(abs(IPD))	0.058 ms	0.05 ms	0.1 ms	0.19 ms	0.1 ms	0.19 ms	1 ms	0 ms
IQR(abs(IPD))	0.048 ms	0.042 ms	0.09 ms	0.11 ms	0.11 ms	0.2 ms	0.009 ms	0.01 ms
Loss	0%	0%	0%	0.008%	0%	0%	0%	0%

Since the timestamps of measurements for one MA to a target are not synchronized with another MA to the same target, they are sampling the network at different times. Thus we decided not to use the residuals in the RTTs between one pair and another. Typically the difference in the time of a measurement from say pinger.slac.stanford.edu to sitka.triumf.ca versus pinger-raspberry.slac.stanford.edu to sitka.triumf.ca averages at 8 mins (see spreadsheet mtr-localhost tab diff mtr).

To find the probability of the distributions overlapping we can use a nomogram of mean differences versus error ratios given in **Overlapping Normal Distributions**. John M. Linacre for normal distributions. However this does not cover the range we are interested in.

We used the Z Tests to compare two samples, See for example "Comparing distributions: Z test" available at <http://homework.uoregon.edu/pub/class/es202/ztest.html>

Spreadsheet (ipdv-errors) has min/avg/max, % and loss between and from pinger & pi to & from sitka, and to cern, plus histograms probabilities, IPDV errors.

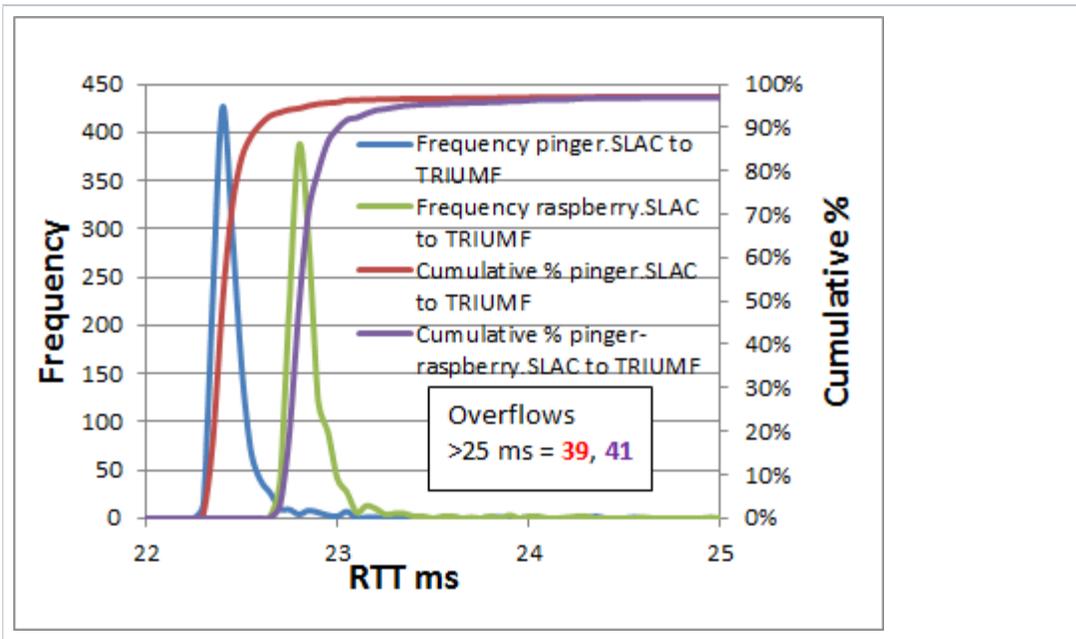
Spreadsheet (prob) (Probability table as submitted, IPDV from SLAC to TRIUMF. Also images of IPD graphs from frequency.pl. The IPD distribution from pinger to SLAC has the figure below.

However the ping distributions are decidedly non-normal (see for example the figure below) have wide outliers, and are heavy tailed on the upper side (see <https://www.slac.stanford.edu/comp/net/wan-mon/ping-hi-stat.html>). This leads to large standard deviations (one to two order of magnitude greater than the IQR) in the RTT values. As can be seen from the table this results in low values of the Z-test and a false probability of no significant statistical difference. Using the IQRs of the frequency distributions instead generally leads to much higher values of the Z-test and hence a higher probability that the distributions of RTTs between two pairs of hosts are significantly different. Comparing the frequency distributions it is seen that there is indeed a marked offset in the RTT values of the peak frequencies and a resultant difference in the cumulative RTT distributions, Using the non-parametric Kolomogorov Smirnof test (KS test) also indicated significant differences in the distributions.

Average RTTs from pinger.SLAC and pinger-raspberry.slac to TRIUMF

(Figure can be found in <https://confluence.slac.stanford.edu/download/attachments/192190786/ipdv-errors.xls>)

The plot cuts off at about the cumulative 98%. The number of overflows not plotted are indicated.

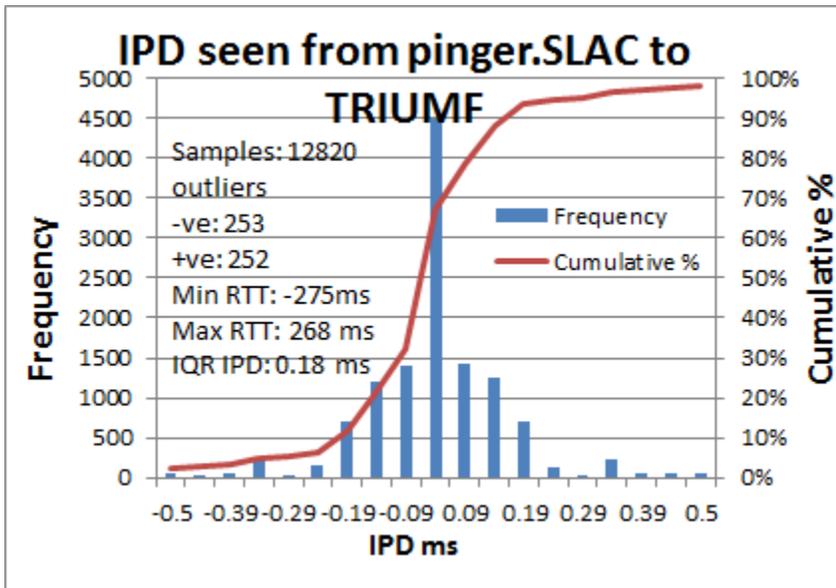


The RTT measurements made from pinger.SLAC and pinger-raspberry.SLAC to TRIUMF and CERN average around 23ms and 151ms respectively. Despite this large difference in average RTTs, comparing the average RTTs from pinger.SLAC with those from pinger-raspberry.SLAC yields a difference of only ~0.35ms for both TRIUMF and CERN.

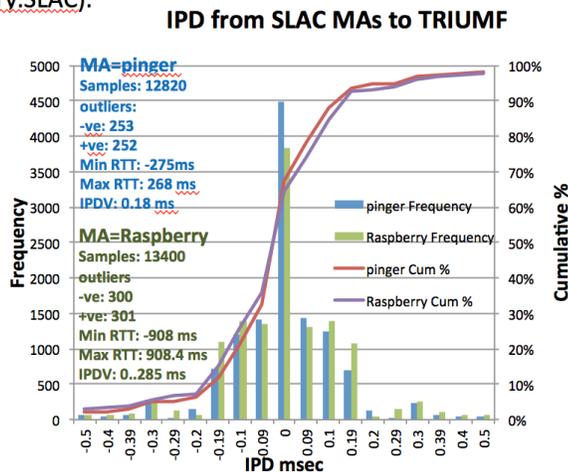
Looking at the traceroutes, using Matt's traceroute to measure the RTT to each hop, indicates that this difference starts at the first hop and persists for later hops. See spreadsheet (MTR tab). We therefore made ping measurements from each SLAC MA to its loopback network interface. The measurements were made at the same times to facilitate comparisons. They indicate that the pinger-raspberry.SLAC is ~0.13ms slower in responding than the pinger.SLAC MA. Thus approximately 1/3rd of the difference in average RTS to TRIUMF measured by the two SLAC MAs is due to the MA platform itself.

IPDV

PingER's main metric for measuring jitter is the IPDV. A typical IPD distribution from which the IPD is derived is shown below.



Plot of the IPD distributions for pinger.SLAC and pinger-raspberry.SLAC to TRIUMF. The plot is cut off at IPD = +/- 0.5 ms (approximately below the 2% threshold and above the 98% threshold). The number of samples and outliers are shown as are the minimum and maximum values of IPD, and the IPDV (blue for pinger.SLAC and green for pinger-raspberry.SLAC).



IPD distributions are centered on 0ms and have very wide tails. The one in the figure is cut off below the 2 percentile and above the 98% percentile. The number of outliers not shown is given in the figure, as are the maximum and minimum values of IPD. The distribution is thus seen to have very positive and negative tails. Also as illustrated in the figure a typical IPD distribution has a very sharp peak. To derive the IPDV we take the IQR of the IPD distribution. The values for the IPDV for the various measurements are shown in the following table; The errors(S) in the IPD are taken from the IQR for the hourly PingER IPDVs observed for the same period. It is seen that the Z-Test in this indicates a value of < 2.0. Assuming the Z-Test is relevant for the non-normal IPD distributions if one uses the IQRs instead of the standard deviation, a value of < 2 for the Z-test statistics indicates the two samples are the same (see <http://homework.uoregon.edu/pub/class/es202/ztest.html>).

Powerpoint of figures.

Localhost comparisons

To try and determine more closely the impact of the host on the measurements, we eliminated the effect of the network by making measurements to the host's localhost NIC port. Looking at the analyzed data it is apparent that:

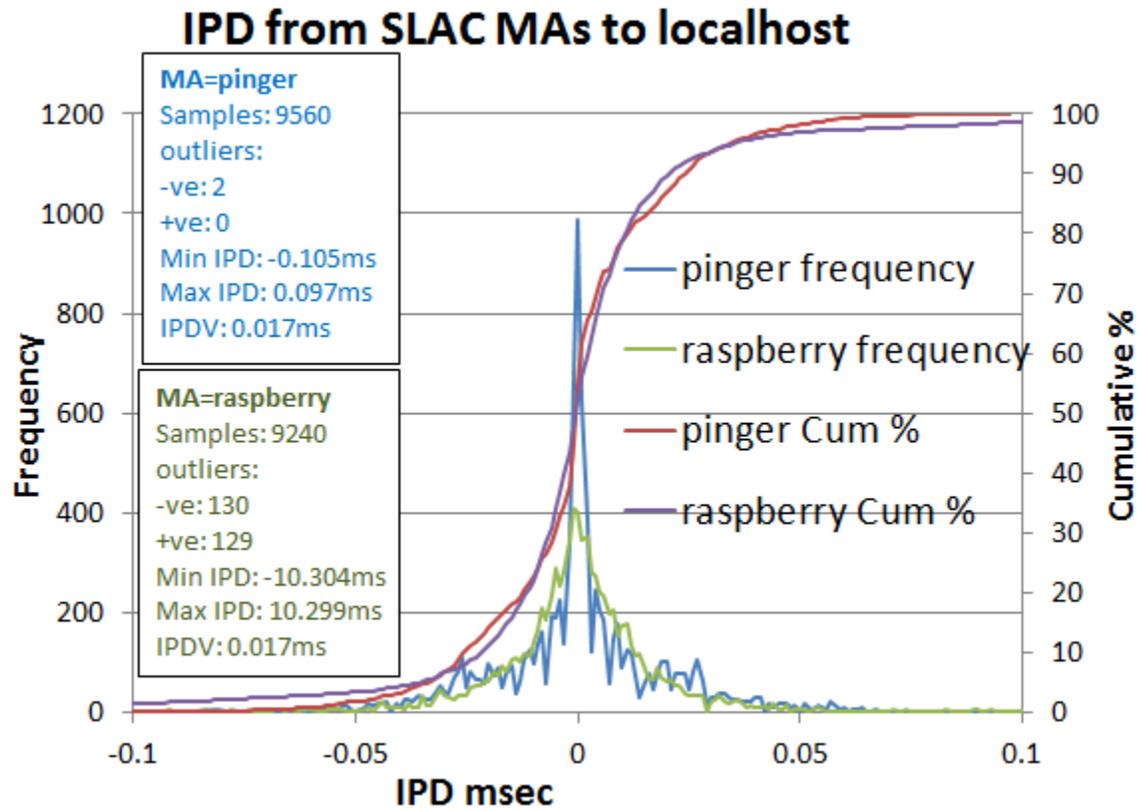
- the pinger hosts has considerably shorter RTTs

- the pinger host has fewer large RTT outliers

Looking at the RTT frequency distributions it is apparent that:

- the pinger and raspberry host RTT frequency distributions hardly overlap at all (one sample (0.04%) of the raspberry pi distribution overlaps ~ 7% of the pinger distribution).
- the pinger host's Median RTT is 7 times smaller (0.03ms vs 0.2 ms) than that of the raspberry-pi
- the pinger host's maximum outlier (0.174ms) is ~ factor 60 smaller than the raspberry-pi.
- the pinger host exhibits a pronounced bimodality not seen for the raspberry-pi.
- the IQRs are very similar (0.025 vs 0.020ms)

Looking at the IPD distributions it is apparent that:



- the pinger distribution is much sharper (max frequency 987 vs 408 for raspberry-pi)
- there is a slight hint of multi-modality in the pinger distribution less visible in the raspberry-pi distribution
- the minimum and maximum outliers are about a factor 100 lower for pinger than the raspberry-pi (~ +/-0.1 vs ~10ms)
- the IPDV's are similar (0.017ms), whereas the standard deviations differ by almost a factor of 10 (0.021ms vs. 0.227ms). This further illustrates the impact of smaller outliers seen by pinger.

Comparing the IPD frequency distributions for the pinger.SLAC and pinger-raspberry MAs to their localhosts with those to TRIUMF it is seen that the differences are washed out as one goes to the longer RTTs.

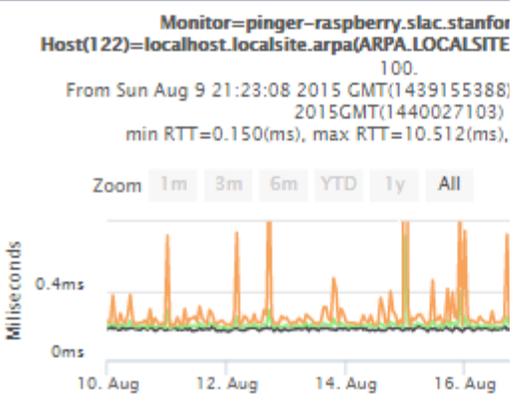
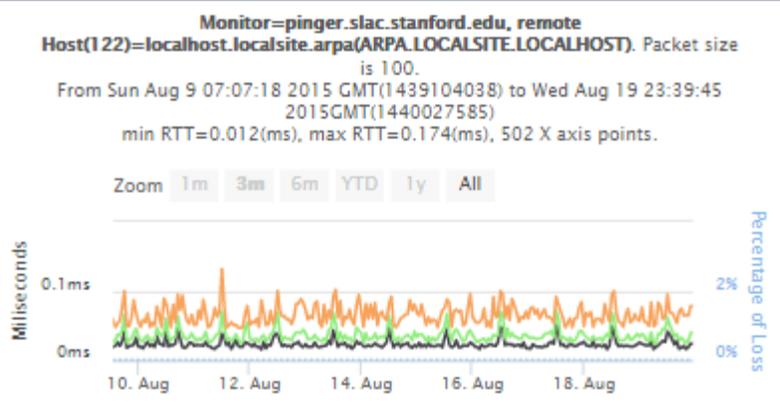
Though there are large relative differences in the distributions, the absolute differences of the aggregated statistics (medians, IQRs) are sub millisecond and so should not noticeably affect PingER wide area network results.

Though we have not analyzed the exact reason for the relative differences, we make the following observations:

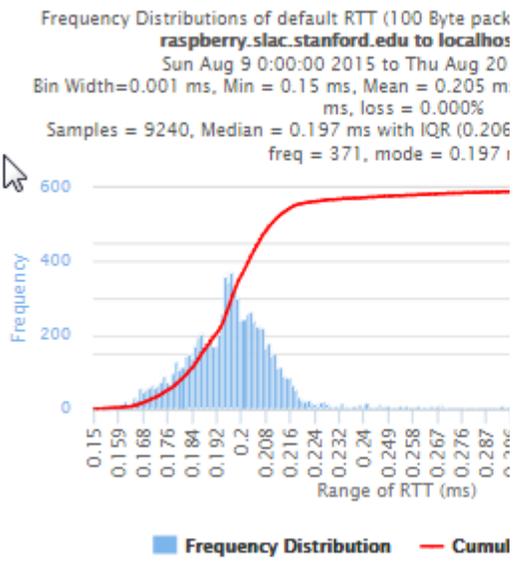
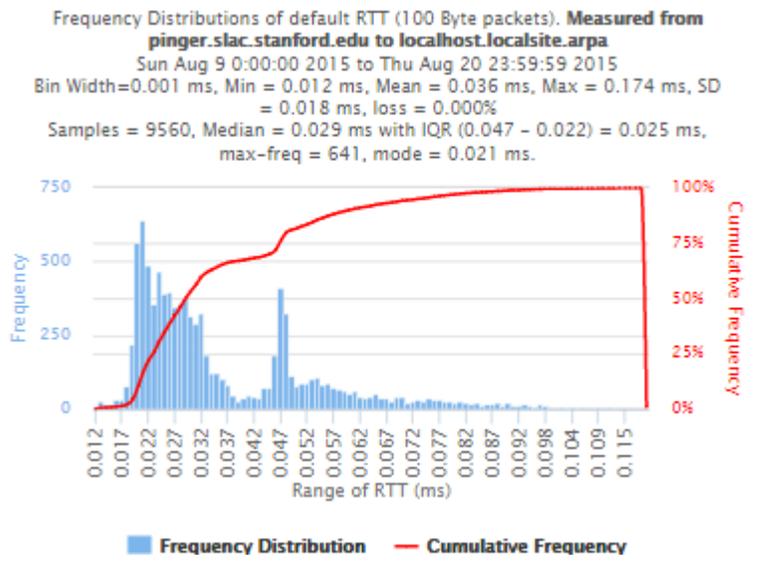
- The factor of 7 difference in the median RTTs is probably at least partially related to the factor ten in the NICs (1Gbps vs 100Mbps) and the factor of 5 difference in the clock speeds (3Ghz vs 600MHz).
- Looking at the minimum RTTs for each set of pings it is seen that pinger has a larger ratio of median RTT / IQR compared to the raspberry-pi (0.33 vs 0.096).
 - This may be related to the fact that the pinger host is less dedicated to acting as a PingER MA since in addition it runs lots of cronjobs to gather, archive and analyze the data. This may also account for some of the more pronounced multi-modality of the pinger host's RTT distributions.
- Currently we do not have a rationale for the reduced outliers on pinger vs the raspberry-pi.

pinger.slac.stanford.edu	pinger-raspberry.slac.stanford.edu
--------------------------	------------------------------------

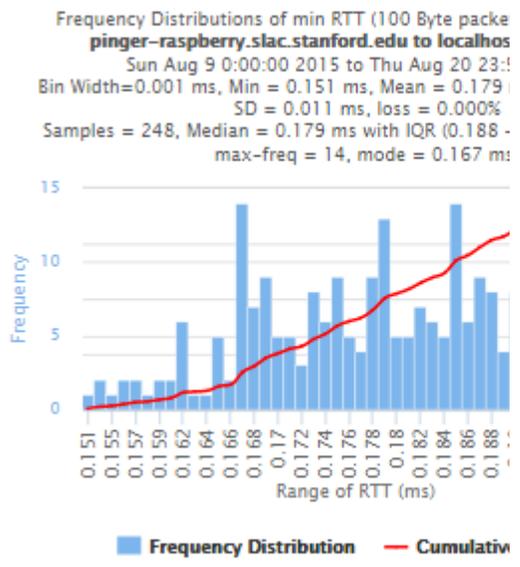
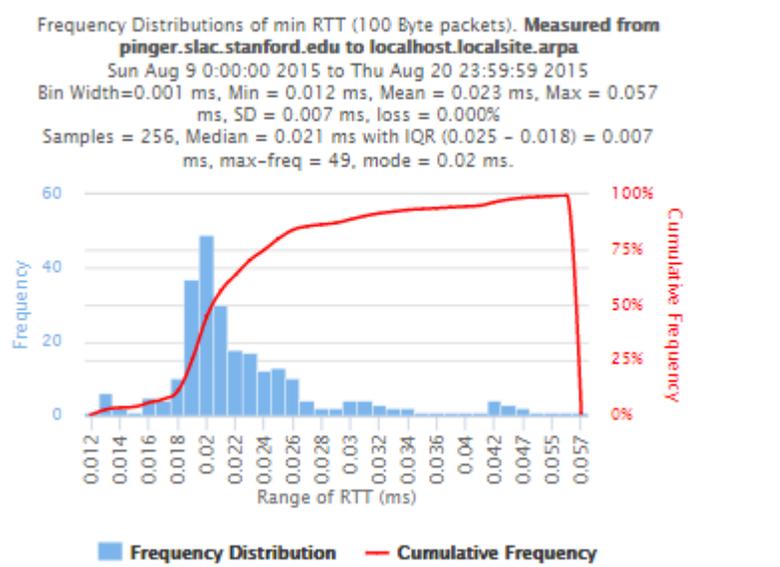
Time series

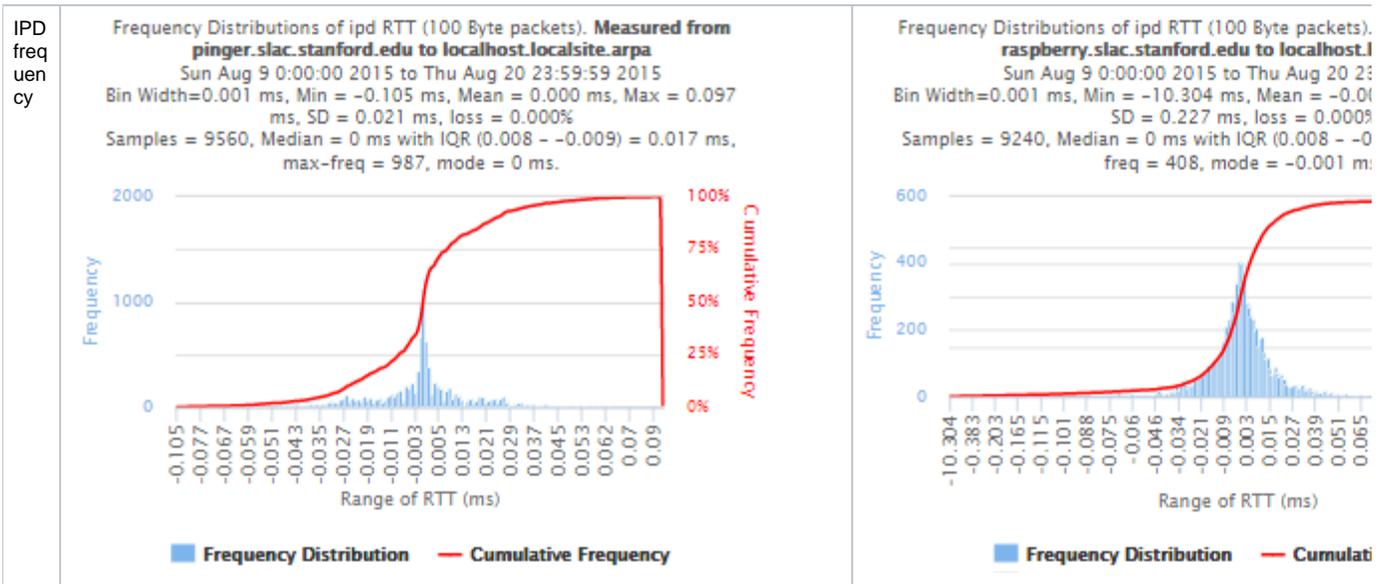


RTT frequency



Minimum RTTs for each set of 10 pings





Kolmogorov-Smirnov Test

The Kolmogorov-Smirnov test (KS-test) tries to determine if two datasets differ significantly. The KS-test has the advantage of making no assumption about the distribution of data. In other words it is non-parametric and distribution free. The method is explained here and makes use of an Excel tool called "Real Statistics". The tests were made using the raw data and distributions, both methods had similar results except for the 100Bytes Packet that had a great difference in the results. The results using raw data says both samples does not come from the same distribution with a significant difference, however if we use distributions the result says that only the 1000Bytes packet does not come from the same distribution. Below you will find the graphs for the distributions that were created and the cumulative frequency in both cases plotted one above other (in order to see the difference between the distributions).

	Raw data - 100 Packets	Distribution - 100 Packets	Raw data - 1000 Packets	Distribution - 1000 Packets
D-stat	0.194674	0.039323	0.205525	0.194379
P-value	4.57E-14	0.551089	2.07E-14	7.32E-14
D-crit	0.0667	0.067051	0.0667	0.067051
Size of Raspberry	816	816	816	816
Size of Pinger	822	822	822	822
Alpha	0.05i			

If D-stat is greater than D-crit the samples are not considered from the same distribution with a (1-Alpha) of accuracy. Remember that D-stat is the maximum difference between the two cumulative frequency curves.

$$D_{m,n} = \max_x |F(x) - G(x)|$$

Source: <http://www.real-statistics.com/non-parametric-tests/two-sample-kolmogorov-smirnov-test/>

In these links you can find the files containing the graphs, histograms and the complete analysis for Kolmogorov-Smirnov between Pinger and Pinger-Raspberry and for SITKA.

Unknown macro: 'html'